

Hebbian Crosstalk and Input Segregation

Anca Rădulescu¹, Paul Adams^{2,3}

August 2, 2012

¹Department of Mathematics, 395 UCB, University of Colorado, Boulder, *radulesc@colorado.edu*

²Department of Neurobiology and Behavior, Stony Brook University, Stony Brook, *padams@notes.cc.sunysb.edu*

³Kalypso Institute, Stony Brook, NY 11790

Abstract

Purpose. We previously proposed that Hebbian adjustments that are incompletely synapse specific (“crosstalk”) might be analogous to genetic mutations. We analyze aspects of the effect of crosstalk in Hebbian learning using the classical Oja model.

Methods. In previous work we showed that crosstalk leads to learning of the principal eigenvector of **EC** (the input covariance matrix pre-multiplied by an error matrix that describes the crosstalk pattern), and found that with positive input correlations increasing crosstalk smoothly degrades performance. However, the Oja model requires negative input correlations to account for biological ocular segregation. Although this assumption is biologically somewhat implausible, it captures features that are seen in more complex models. Here, we analyze how crosstalk would affect such segregation.

Results. We show that for statistically unbiased inputs crosstalk induces a bifurcation from segregating to non-segregating outcomes at a critical value which depends on correlations. We also investigate the behavior in the vicinity of this critical state and for weakly biased inputs.

Conclusions. Our results show that crosstalk can induce a bifurcation under special conditions even in the simplest Hebbian models and that even the low levels of crosstalk observed in the brain could prevent normal development. However, during learning pairwise input statistics are more complex and crosstalk-induced bifurcations may not occur in the Oja model. Such bifurcations would be analogous to “error catastrophes” in genetic models, and we argue that they are usually absent for simple linear Hebbian learning because such learning is only driven by pairwise correlations.

Keywords. Crosstalk, Hebbian synapses, pairwise correlations, sensitivity analysis, codimension two bifurcation.

1 Introduction

1.1 Background

Learning is thought to occur as a result of changes in synaptic strength triggered by pre- and postsynaptic neural activity, in a “Hebbian” manner. Such changes are not completely specific to the synapses at which the activity occurs [23, 7, 9], because of inevitable albeit minimal second-messenger diffusion.

Oja [33] showed that a simple model neuron could perform unsupervised Hebbian learning of the first principal component of an input distribution. In this model, unlimited weight growth is prevented using an additional term in the learning rule, producing an implicit, “multiplicative” weight normalization [30]. Biological synapses do show Hebbian properties, using well-understood, spike-coincidence detection machinery, raising the possibility that real neurons can exhibit similar unsupervised learning. Finding principal components could be very useful in the brain for data compression and transmission, since for Gaussian data such representations have statistically optimal properties, and often neural signals are approximately Gaussian. Furthermore, representational learning often requires that inputs be pairwise decorrelated. Hebbian learning can also explain developmental changes, such as the segregation of visual input to central neurons.

Recent data suggest [23, 7, 9] that weight updates may be affected by each other, for example due to unavoidable residual second messenger diffusion between closely spaced synapses. We have suggested that such crosstalk is analogous to mutation in genetics, and that cortical circuitry may be specialized to reduce it. However, it is not clear that learning would be subject to an “error catastrophe” such as that occurring in genetic systems [14]. If complete learning failure does not occur at a critical, low, crosstalk level, such circuitry might not be necessary.

In a recent paper [36] we examined how crosstalk would affect the Oja model. We considered a learning network consisting of a single output neuron receiving, through a set of n input neurons, n signals $\mathbf{x} = (x_1, \dots, x_n)^T$ drawn from a probability distribution $\mathcal{P}(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^n$, transmitted via synaptic connections of strengths $\boldsymbol{\omega} = (\omega_1, \dots, \omega_n)^T$. The resulting scalar output y was generated as the weighted sum of the inputs $y = \mathbf{x}^T \boldsymbol{\omega}$.

The synaptic weights ω_i were modified in accordance with Oja’s rule of learning, by implementing first a Hebb-like strengthening proportionally with the product of x_i and y (with small constant of proportionality, or *learning rate*, γ)

$$\omega_i(t+1) = \omega_i(t) + \gamma y(t) x_i(t)$$

followed by an approximate “normalization” step (applicable for small γ and $\|\mathbf{w}\|$ close to one), maintaining the Euclidean norm of the weight vector $\boldsymbol{\omega} = (\omega_1, \dots, \omega_n)^T$ close to one.

$$\boldsymbol{\omega}(t+1) = \boldsymbol{\omega}(t) + \gamma y(t) [\mathbf{x}(t) - y(t) \boldsymbol{\omega}(t)]$$

We considered the long-term average of this Oja equation, using the input covariance matrix $\mathbf{C} = \mathbf{x}^T \mathbf{x}$ as an appropriate long-term characterization of the inputs, and studying the behavior of $\mathbf{w}(t) = \langle \boldsymbol{\omega}(t+1) | \boldsymbol{\omega}(t) \rangle$:

$$\frac{d\mathbf{w}}{dt} = \gamma [\mathbf{C}\mathbf{w} - (\mathbf{w}^T \mathbf{C}\mathbf{w}) \mathbf{w}]$$

in continuous time, or:

$$\Delta \mathbf{w} = \gamma [\mathbf{C}\mathbf{w} - (\mathbf{w}^T \mathbf{C}\mathbf{w}) \mathbf{w}]$$

as a discrete time approximation.

We then introduced inspecificity into the learning equation [36]. We implemented this inspecificity by assuming that, on average, only a fraction q of the intended update reaches the appropriate connection, the remaining fraction $1 - q$ being distributed amongst the other connections (according to a rule which we defined according to plausible underlying biology). The quality factor q is analogous to a similar factor in molecular evolution theory that represents the fidelity of single-base copying [41]. The actual update at a given connection thus includes contributions from erroneous or inaccurate updates from other connections. The erroneous updating process was formally described by an error matrix $\boldsymbol{\mathcal{E}}$, independent of the inputs, whose elements, which depend on average on q , reflect at each time step t the fractional contribution that the activity through the connection with weight ω_i makes to the update of ω_j .

$$\omega_i(t+1) = \omega_i + \gamma y([\boldsymbol{\mathcal{E}}\mathbf{x}]_i - y\omega_i)$$

The discrete long-term statistics can be then written in matrix form as:

$$\Delta \mathbf{w} = \gamma [\mathbf{E} \mathbf{C} \mathbf{w} - (\mathbf{w}^T \mathbf{C} \mathbf{w}) \mathbf{w}]$$

where the “error matrix” $\mathbf{E} = \langle \mathcal{E} \rangle$ is a symmetric matrix with positive entries, which equals the identity matrix $\mathbf{I} \in \mathcal{M}_n(\mathbb{R})$ in case of perfect quality updates. Then the rule changes into:

$$\frac{d\mathbf{w}}{dt} = \gamma [\mathbf{E} \mathbf{C} \mathbf{w} - (\mathbf{w}^T \mathbf{C} \mathbf{w}) \mathbf{w}] \quad (1)$$

Throughout the paper, we call this the (inspecific) Oja rule with continuous time updates.

We studied the asymptotic behavior of this n-dimensional system, starting with a local linear analysis of the equilibria and their stability. Although this rule is nonlinear, the Hebbian update term is linear in the output, and we sometimes refer to this, and related, rules, as being “linear,” in contrast to other Hebbian rules [25, 24, 6, 34, 15, 21, 10] which are nonlinear in the output.

Note that the symmetric, positive definite matrix $\mathbf{C} \in \mathcal{M}_n(\mathbb{R})$ defines a dot product between any two vectors \mathbf{w} and \mathbf{v} in \mathbb{R}^n as $\langle \mathbf{v}, \mathbf{w} \rangle_{\mathbf{C}} = \mathbf{v}^T \mathbf{C} \mathbf{w}$. Although both \mathbf{C} and \mathbf{E} are symmetric, the product $\mathbf{E} \mathbf{C}$ is not symmetric in the Euclidean metric. However, in a new metric defined by the dot product $\langle \cdot, \cdot \rangle_{\mathbf{C}}$, $\mathbf{E} \mathbf{C}$ is symmetric: $\langle \mathbf{E} \mathbf{C} \mathbf{u}, \mathbf{v} \rangle_{\mathbf{C}} = (\mathbf{E} \mathbf{C} \mathbf{u})^T \mathbf{C} \mathbf{v} = \mathbf{u}^T \mathbf{C}^t \mathbf{E}^t \mathbf{C} \mathbf{v} = \mathbf{u}^T \mathbf{C} \mathbf{E} \mathbf{C} \mathbf{v} = \langle \mathbf{u}, \mathbf{E} \mathbf{C} \mathbf{v} \rangle_{\mathbf{C}}$, for all $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$. Hence $\mathbf{E} \mathbf{C}$ has a basis of eigenvectors, orthogonal with respect to the dot product $\langle \cdot, \cdot \rangle_{\mathbf{C}}$. The following is immediate:

Description of equilibria of the system (1). *An equilibrium for the system is any vector $\mathbf{w} = (w_1 \dots w_n)^T$ such that $\mathbf{E} \mathbf{C} \mathbf{w} = (\mathbf{w}^T \mathbf{C} \mathbf{w}) \mathbf{w}$, i.e., an eigenvector of $\mathbf{E} \mathbf{C}$ (with corresponding eigenvalue $\lambda_{\mathbf{w}}$), normalized, with respect to the norm $\|\cdot\|_{\mathbf{C}} = \langle \cdot, \cdot \rangle_{\mathbf{C}}$, so that $\|\mathbf{w}\|_{\mathbf{C}} = \lambda_{\mathbf{w}}$. Generically, $\mathbf{E} \mathbf{C}$ has a strictly positive, unique maximal eigenvalue, and the corresponding eigendirection is orthogonal in $\langle \cdot, \cdot \rangle_{\mathbf{C}}$ to all other eigenvectors of $\mathbf{E} \mathbf{C}$.*

For an equilibrium \mathbf{w} of the system (1), the Jacobian matrix $Df_{\mathbf{w}}^{\mathbf{E}}$ around \mathbf{w} is (see Appendix 1):

$$Df_{\mathbf{w}}^{\mathbf{E}} = \gamma [\mathbf{E} \mathbf{C} - 2\mathbf{w}(\mathbf{C} \mathbf{w})^T - (\mathbf{w}^T \mathbf{C} \mathbf{w}) \mathbf{I}] \quad (2)$$

Then we have the following (see Appendix 1 for proof):

Stability criteria for equilibria. *Suppose $\mathbf{E} \mathbf{C}$ has a multiplicity one largest eigenvalue. A normalized eigenvector \mathbf{w} is a local hyperbolic attracting equilibrium for (1) iff it corresponds to the maximal eigenvalue of $\mathbf{E} \mathbf{C}$.*

Such attractors always exist provided $\mathbf{E} \mathbf{C}$ has a maximal eigenvalue of multiplicity one, which is generically true. Then the network learns, depending on its initial state, one of the two stable equilibria, which are the two (opposite) maximal eigenvectors of the modified input distribution, normalized so that $\|\mathbf{w}\|_{\mathbf{C}} = \lambda_{\mathbf{w}}$. It can be shown easily that these two attractors (the appropriately normalized eigenvectors corresponding to the maximal eigenvalue of $\mathbf{E} \mathbf{C}$) can be the only attractors in the system (see Appendix 2 for proof).

In a previous paper [36], we further analyzed the sensitivity of the system under variations of parameters, for some biologically plausible forms of the covariance and error matrices:

$$\mathbf{C} = \begin{bmatrix} v + \delta_1 & c & \cdots & c \\ c & v + \delta_2 & \cdots & c \\ \vdots & & \ddots & \vdots \\ c & c & \cdots & v + \delta_n \end{bmatrix} \text{ and } \mathbf{E} = \begin{bmatrix} q & \epsilon & \cdots & \epsilon \\ \epsilon & q & \cdots & \epsilon \\ \vdots & & \ddots & \vdots \\ \epsilon & \epsilon & \cdots & q \end{bmatrix}$$

where the input covariance matrix had uniform covariances $c > 0$ and variance biases $\delta_1 \geq \delta_2 \geq \dots \geq \delta_n$; the error matrix was defined such that $q > \epsilon > 0$, $q + (n - 1)\epsilon = 1$. Our analysis of this system concluded that the effect of biologically realistic levels of crosstalk would typically only produce small gradual changes in the learning process, though when inputs carry very similar signals, the effects could be more dramatic. In this paper we explore this “very similar” scenario more thoroughly. In particular we describe the effect of crosstalk in the special “unbiased” case, where the inputs have identical statistics.

1.2 Biased and unbiased inputs

Our previous analysis considered only distributions of inputs with a bias in the covariance matrix (we imposed the condition that $\mathbf{E} \mathbf{C}$ has a leading eigenvalue of multiplicity one). While this case is mathematically generic, previous

work using related models (without crosstalk [31]) to study learning in the visual system, often assumed that the input statistics are “unbiased,” or identical for each input (for example, because inputs from corresponding points in the left and right eyes see the same point in visual space). It is well known that in the two-dimensional case, if the two inputs x_1 and x_2 are positively correlated (as one might anticipate for active vision), linear Hebbian learning does not predict the observed developmental segregation of visual afferents [13, 10, 42, 45]; negative correlations (or a nonlinear rule) are required. However, modifications in learning rules, for example subtractive normalization [22, 45, 32, 27], a weight-dependent rule [18] or a BCM rule [10], although not always originally developed to explain segregation, can overcome this difficulty. Often these rules lead to Hebbian learning driven by a modified version of the covariance matrix. In the current work, we examine the dynamics of Oja learning with crosstalk when inputs are unbiased, and how this changes when a slight bias is introduced.

We show here that in the unbiased negative correlation case, the system undergoes a bifurcation in dynamics at a critical crosstalk level. Related results have been obtained by [17, 11]. While there is no true bifurcation in the near-unbiased case, the very dramatic change in learning that occurs over a small error range is biologically indistinguishable from a true bifurcation. We discuss our results in relation to models of development and learning.

1.3 Our current model

In this paper, we will consider the continuous-time, two-dimensional nonlinear rule of Oja (i.e., for two input channels and one output), with covariance matrix \mathbf{C} and error matrix \mathbf{E} symmetric matrices having the forms $\mathbf{C} = \begin{pmatrix} v + \delta & c \\ c & v \end{pmatrix}$ and $\mathbf{E} = \begin{pmatrix} q & 1 - q \\ 1 - q & q \end{pmatrix}$. The parameters are such that $1/2 < q \leq 1$, $v > 0$ and $c < 0$, such that $v > |c|$, $v > |\delta|$ and $v(v + \delta) > c^2$ (i.e., $\det(\mathbf{C}) > 0$). The 2D system expands to:

$$\begin{aligned} \dot{w}_1 &= [q(v + \delta) + (1 - q)c]w_1 + [qc + (1 - q)v]w_2 - [vw_1^2 + 2cw_1w_2 + vw_2^2]w_1 \\ \dot{w}_2 &= [(1 - q)(v + \delta) + qc]w_1 + [(1 - q)c + qv]w_2 - [vw_1^2 + 2cw_1w_2 + vw_2^2]w_2 \end{aligned} \quad (3)$$

The rest of the paper is centered around this 2-dimensional model. In Section 2, we establish the mathematical background of the model’s behavior. We analyze some of its local and global dynamics, observe the dependence of these dynamics on parameters and discuss bifurcations. One of the phenomena central to our interest is how the behavior of the system changes when the bias parameter δ varies, in particular when it approaches zero (i.e., the inputs are very close to a perfectly unbiased state). In Section 3, we discuss the results in the context of visual modeling and ocular segregation of inputs.

2 Linear analysis of the 2D dynamics

We notice that the phase plane of the system is symmetric about the origin (i.e., if $w(t)$ is a solution curve for the system, then $-w(t)$ is as well). The trace, determinant and eigenvalues of \mathbf{EC} can be obtained easily as expressions of the system parameters:

$$\det(\mathbf{EC}) = \det(\mathbf{E}) \det(\mathbf{C}) = (2q - 1)[v(v + \delta) - c^2] > 0$$

$$\text{tr}(\mathbf{EC}) = 2(1 - q)c + q(2v + \delta) > 0 \text{ (from the Cauchy-Schwartz inequality).}$$

Lemma 2.1. *For all parameter values, \mathbf{EC} has two real eigenvalues $\mu_{1,2}$, which are distinct unless the conditions $\delta = 0$ and $q = q^* = \frac{v}{v - c}$ are simultaneously satisfied. More precisely, when $\mu_1 \neq \mu_2$, we have:*

$$\begin{aligned} \mu_1 &= \frac{2(1 - q)c + q(2v + \delta) + \sqrt{\Delta}}{2} \text{ larger eigenvalue, with eigenline of slope } z_1 = \frac{-q\delta + \sqrt{\Delta}}{2\beta} \\ \mu_2 &= \frac{2(1 - q)c + q(2v + \delta) - \sqrt{\Delta}}{2} \text{ smaller eigenvalue, with eigenline of slope } z_2 = \frac{-q\delta - \sqrt{\Delta}}{2\beta} \end{aligned}$$

where $\beta = qc + (1 - q)v$ and $\Delta = [2qc + (1 - q)(2v + \delta)]^2 + (2q - 1)\delta^2$.

Proof. The calculation of eigenvalues and eigenvectors is immediate from the characteristic equation of \mathbf{EC} : $X^2 - \text{tr}(\mathbf{EC})X + \det(\mathbf{EC}) = 0$, with discriminant

$$\Delta = \text{tr}(\mathbf{EC})^2 - 4\det(\mathbf{EC}) = [2qc + (1-q)(2v+\delta)]^2 + (2q-1)\delta^2$$

Notice that $\Delta \geq 0$, with equality $\Delta = 0$ (i.e., double eigenvalue for \mathbf{EC}) iff both $\delta = 0$ and $\beta = 0$. The critical quality value (where the two eigenvalues are equal, producing a switch in the dynamics when $\delta = 0$) is $q^* = \frac{v}{v-c}$. Since $v > |c|$, this value occurs within the appropriate q range, $(1/2, 1]$ (see Figure 1). \square




Figure 1: **Evolution of the eigenvalues as the quality q is varied, in three different δ slices.** $\delta = -0.2$ (A), $\delta = 0$ (B) and $\delta = 0.5$ (C). Fixed parameters: $v = 1$ and $c = -0.4$, hence $q^* = 1/1.4 \sim 0.71$. When $\delta = 0$, the eigenvalues μ_1 and μ_2 touch at $q = q^*$. For $\delta \neq 0$, the two curves avoid this crossing; the minimal distance between them occurs at $q = q^*$, but it is strictly positive.

2.1 Equilibria of the 2D system

Throughout this section, in addition to working in our generally specified parameter ranges, we will assume that \mathbf{EC} has distinct eigenvalues (i.e., $\delta \neq 0$ or $q \neq q^*$). In this case, the system has as equilibria the origin $\mathbf{w} = 0$, and two pairs of opposite eigenvectors of \mathbf{EC} normalized such that $\mathbf{w}^t \mathbf{C} \mathbf{w} = \mu$ (where μ is the respective eigenvalues of each pair).

The normalization condition can be written as:

$$\mathbf{w}^t \mathbf{C} \mathbf{w} = (v + \delta)w_1^2 + 2cw_1w_2 + vw_2^2 = \mu$$

Using the same notation $z = w_2/w_1$, this can be rewritten as $vz^2 + 2cz + (v + \delta) = \mu/w_1^2$, so that

$$\|\mathbf{w}\| = \sqrt{\frac{\mu(z^2 + 1)}{vz^2 + 2cz + (v + \delta)}}$$

Thus the norm varies with both error and correlation. The position and stability of the four nonzero equilibria vary with the parameters v, c, δ and q . If we aim to study the sensitivity of the system's dynamics under parameter perturbations, the next step should be establishing the linear stability of these equilibria; this follows directly from the general results in Section 1:

Description and stability of equilibria. *Suppose the matrix \mathbf{EC} has distinct eigenvalues. The system (3) has five distinct equilibria, $\mathbf{w} = 0$ and four normalized eigenvectors of \mathbf{EC} . The two (opposite) eigenvectors of the larger eigenvalue are hyperbolic attractors, and the two (opposite) eigenvectors corresponding to the lower eigenvalue are saddles. The origin is repelling.*

More precisely, this means that if $\mu_{\mathbf{w}}$ is the larger eigenvalue of \mathbf{EC} , the Jacobian matrix $D_{\mathbf{w}}$ has two negative eigenvalues, hence \mathbf{w} is an attracting node. If instead $\mu_{\mathbf{w}}$ is the smaller eigenvalue of \mathbf{EC} , then $D_{\mathbf{w}}$ has two real eigenvalues of opposite signs, and \mathbf{w} is a saddle equilibrium.

We are particularly interested in the behavior near and at $\delta = 0$. The above characterization of equilibria applies when $\delta \neq 0$, but it breaks down in the parameter slice $\delta = 0$, at the critical point when \mathbf{EC} has a double eigenvalue. In other words we expect that the system undergoes a bifurcation in the unbiased $\delta = 0$ slice, which does not exist

in the other, $\delta \neq 0$ slices (i.e., when “bias” is present in the inputs), therefore we will study this case separately.

For the following paragraph (Section 2.2) we assume $\delta \neq 0$. The unbiased case $\delta = 0$ is discussed separately in Section 2.3. The results are integrated and concluded in Section 2.4.

2.2 More properties of the phase plane

One way to describe the dynamics of the system, including the more global aspects and possibly cyclic behavior (which has not yet been excluded) is to follow the rotational direction of the solution trajectories in different regions of the (w_1, w_2) phase-plane under the velocity field (\dot{w}_1, \dot{w}_2) .

Consider the angle $\theta \in [-\pi/2, \pi/2]$ made by the direction (w_1, w_2) with the w_1 axis. As before, call $z = w_2/w_1 = \tan(\theta)$ and $\beta = qc + (1 - q)v$. Then, along a trajectory in the (w_1, w_2) plane,

$$\begin{aligned}\dot{z} &= \frac{d}{dt} \left(\frac{w_2}{w_1} \right) = \frac{\dot{w}_2 w_1 - \dot{w}_1 w_2}{w_1^2} = [(1 - q)(v + \delta) + qc] - q\delta z - [(1 - q)v + qc]z^2 \\ &= -\beta z^2 - q\delta z + [\beta + (1 - q)\delta]\end{aligned}$$

We first want to establish if there are any values of z for which $\dot{z} = 0$. These are the slopes along which the rotational speed of the trajectories is zero; in other words, they would correspond to invariant lines in the phase-plane.

We consider the quadratic equation: $\dot{z} = -\beta z^2 - q\delta z + [\beta + (1 - q)\delta] = 0$. The discriminant is the same as the one of the characteristic equation of **EC**:

$$\Delta = q^2\delta^2 + 4\beta[\beta + (1 - q)\delta] = [2qc + (1 - q)(2v_\delta)]^2 + (2q - 1)\delta^2$$

The solutions of the quadratic equation will be exactly the slopes of the eigendirections of **EC**:

$$z_{1,2} = \frac{-q\delta \pm \sqrt{\Delta}}{2\beta} \in [-\infty, +\infty]$$

proving the following:

Lemma 2.2. *The eigendirections of the matrix **EC** represent invariant lines under the vector field of system (1).*

We want to better describe the phase-plane behavior *between* the invariant lines $z = z_1$ and $z = z_2$. For any fixed $q \in (1/2, q^*) \cup (q^*, 1]$ (i.e., for $\beta \neq 0$), the rotational speed is given by the sign of the quadratic function $f(z) = -\beta z^2 - q\delta z + [\beta + (1 - q)\delta]$. In principle, we then have two situations:

- i. $q \in (1/2, q^*)$ (i.e., $\beta > 0$). Then $z_1 > z_2$, with $z > 0$ in (z_2, z_1) and $z < 0$ on $(-\infty, z_2) \cup (z_1, \infty)$. The phase plane looks schematically as in Figure 2a.
- ii. $q \in (q^*, 1]$ (i.e., $\beta < 0$). Then $z_1 < z_2$, with $z < 0$ in (z_1, z_2) and $z > 0$ on $(-\infty, z_1) \cup (z_2, \infty)$. The phase plane looks schematically as in Figure 2b.

In other words, all trajectories move asymptotically towards the invariant line $z = z_1$.

Since the behavior of the system seems to a large extent dictated by these invariant lines, we study how the positions of these lines change under variations of the quality parameter q . In other words, we want to study the monotonicity of $z_1 = z_1(q)$ and $z_2 = z_2(q)$. We get the following (for detailed proofs and limit-case behavior $\lim_{q \rightarrow q_\pm} z_{1,2}$, see Appendix 3; for illustrations see Figures 2 and 3:

Proposition 2.3. *If $\delta < 0$, then $\frac{dz_{1,2}}{dq} > 0$ and hence both z_1 and z_2 are increasing as $q \in (1/2, q^*) \cup (q^*, 1]$. In the system's phase plane, this corresponds to a continuous counter-clockwise rotation of the two invariant lines. If $\delta > 0$, then $\frac{dz_{1,2}}{dq} < 0$; hence both $z_{1,2}$ are in this case decreasing as $q \in (1/2, q^*) \cup (q^*, 1]$. In the phase plane, this corresponds to a clockwise rotation of the invariant lines.*

Proposition 2.4. *The angles $\theta_{1,2} \in [-\pi/2, \pi/2]$ between each invariant line and the w_1 abscissa are decreasing with respect to the parameter q in case $\delta > 0$, and are increasing with respect to the parameter q in case $\delta < 0$. Moreover, in both cases, the angular rate of change is finite, at all $q \in (1/2, 1]$.*




Figure 2: **Invariant lines and generic phase plane dynamics.** The invariant lines are marked as $z_1(q)$ and $z_2(q)$. The arrows indicate the rotational direction of the vector field between the two invariant lines. This can be obtained in the right vertical half-plane (where we have defined our angle, $\theta \in [-\pi/2, \pi/2]$), then extended by symmetry in the opposite half-plane. For $q < q^*$ we have $z_1 > z_2$ (A.). As q increases, the two invariant lines rotate: clockwise if $\delta > 0$ and anti-clockwise if $\delta < 0$. At $q = q^*$, one of the invariant lines goes through a vertical stage. For $\delta > 0$, θ_2 jumps from $-\pi/2$ to $\pi/2$, hence z_2 has a vertical asymptote at $q = q^*$, and jumps from $z_2 \rightarrow -\infty$ to $z_2 \rightarrow \infty$. For $\delta < 0$, θ_1 jumps from $\pi/2$ to $-\pi/2$, hence z_1 has a vertical asymptote at $q = q^*$, and jumps from $z_1 \rightarrow \infty$ to $z_1 \rightarrow -\infty$. In consequence, after this critical stage, for $q > q^*$, we have $z_1 < z_2$ (B.). Although the rotation is continuous, either z_1 or z_2 has an infinite discontinuity, due to our definition (mod π) of the angles $\theta_1, 2$.




Figure 3: **Transitions of the phase plane and bifurcation at $q = q^*$, in the slice $\delta = 0$.** A. When $q > q^*$, the stable equilibria are the two vectors of norm $\sqrt{q - 1/2}$ (blue dots) along the invariant line of slope $z_1 = -1$; the saddle equilibria are the two eigenvectors of norm 1 (green dots) along the invariant line of slope $z_2 = 1$. As q decreases from $q = 1$ towards $q = q^*$, the saddles remain unchanged, but the attractors gradually approach the origin (their norm $\sqrt{q - 1/2}$ decreases). B. When $q = q^*$, the system traverses a bifurcation state, characterized by an infinite number (an entire ellipse) of neutrally stable equilibria. This critical state permits the swap of stability between the two invariant lines. C. When $q < q^*$, the stable equilibria are now the two vectors of norm 1 (blue dots) along the invariant line of slope $z_1 = 1$, while the saddle equilibria swapped to the two eigenvectors of norm $\sqrt{q - 1/2}$ (green dots) along the invariant line of slope $z_2 = -1$. As q continues to decrease from $q = q^*$ towards $q = 1/2$, the attractors remain unchanged, and the saddles approach the origin (collapsing into the origin in the limit of $q \rightarrow 1/2$).

2.3 Unbiased case $\delta = 0$

For $\delta = 0$ the computations are simpler; however, as mentioned before, the system has an interesting critical transition which does not appear in the $\delta \neq 0$ slices (occurring from the “touching,” or apparent crossing, of the two eigenvalues at $q = q^*$, as shown in Figure 1).

Proposition 2.5. *Suppose $\delta = 0$. The phase plane of the system depends on the value of q as follows:*

- i. *If $q < q^*$, then $\mu_1 = v + c$ is the larger eigenvalue, with eigendirection $z_1 = 1$ and norm of the corresponding attracting equilibria $\|w\| = 1$. $\mu_2 = (2q - 1)(v - c)$ is the smaller eigenvalue, with eigendirection $z_2 = -1$ and norm of the corresponding saddle equilibria $\|w\| = \sqrt{q - 1/2}$.*
- ii. *If $q > q^*$, then $\mu_1 = (2q - 1)(v - c)$ is the larger eigenvalue, with eigendirection $z_1 = -1$ and norm of the corresponding attracting equilibria $\|w\| = \sqrt{q - 1/2}$. $\mu_2 = v + c$ is the smaller eigenvalue, with eigendirection $z_2 = 1$ and norm of the corresponding saddle equilibria $\|w\| = 1$.*
- iii. *If $q = q^*$, the system contains an infinity of half-stable non-isolated equilibria (each direction will contain two opposite equilibria, describing overall an ellipse of equilibria around the origin).*

Proof. For $\delta = 0$, we have $\dot{z} = -\beta(z^2 - 1)$. The situation $q < q^*$ corresponds to $\beta > 0$, and $q > q^*$ corresponds to $\beta < 0$. Parts i. and ii. follow immediately. For $q = q^*$, $\dot{z} = 0$; all lines through the origin are invariant, and each contains two half-stable equilibria. In the Appendix 4, we show that the locus of these equilibria is an ellipse (see dotted curve in Figure 3), and we describe its axes and foci. \square

Remark 1. For $q = 1$, the attracting equilibria lay along the direction $z = -1$, so that $w_1 + w_2 = 0$. A simple way to quantify how far the stable equilibrium $w = (w_1, w_2)$ degrades from this error-free state as the quality q decreases, we can measure how much the sum $S(q) = |w_1 + w_2|$ deviates from zero, the outcome of perfect learning (Figure 4).

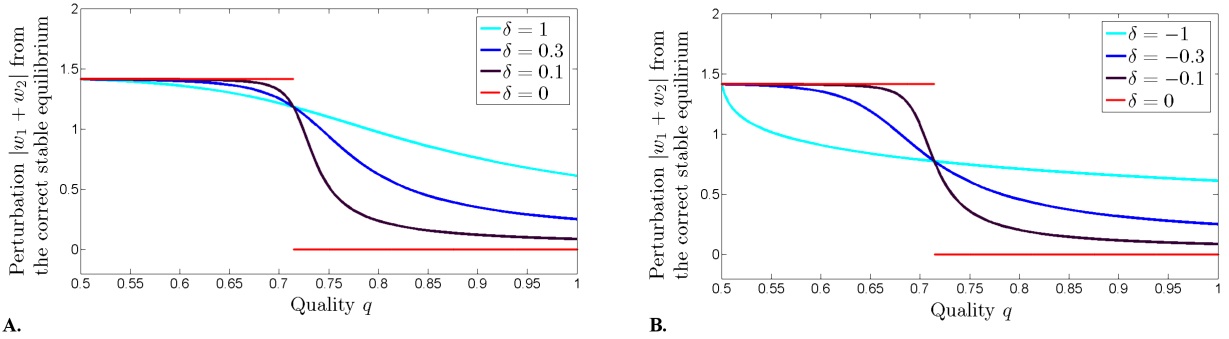


Figure 4: $S(q) = |w_1 + w_2|$ as a measure of the increasing inspecificity of the stable equilibrium, compared to its ideal state $S(1) = 0$, as q decays from $q = 1$. For $v = 1$, $c = -0.4$, we plotted $S(q)$. **A.** For $\delta \geq 0$: $\delta = 1$ (cyan); $\delta = 0.3$ (blue); $\delta = 0.1$ (purple); $\delta = 0$ (red). **B.** For $\delta \leq 0$: $\delta = -1$ (cyan); $\delta = -0.3$ (blue); $\delta = -0.1$ (purple); $\delta = 0$ (red). In both panels, all continuous curves for $\delta \neq 0$ concur at one point, which corresponds to the fact that, for both $\delta > 0$ and $\delta < 0$, the stable equilibrium at $q = q^*$ is independent on the magnitude of δ .

In Figure 1b, the inputs are unbiased ($\delta = 0$), and in the absence of crosstalk ($q = 1$) the inputs segregate completely. As crosstalk increases, the separation between the eigenvalues at first decreases, though the inputs remain completely segregated. However, as crosstalk increases further, the two eigenvalues equalize at the critical quality value $q^* = v/(v - c)$. With further increases in crosstalk, the inputs become completely unsegregated, and the eigenvalues now move apart. This qualitative change at q^* is a bifurcation. Note that although the qualitative behavior only changes at q^* , there is a biologically less important quantitative change: the two symmetric equilibrium weight vectors decrease continuously in length as $\sqrt{q - 1/2}$, as q decreases from $q = 1$ until the bifurcation at $q = q^*$, then remain of unit length for $q < q^*$.

In the slightly unbiased cases $\delta = -0.2$ and 0.5 (Figure 1a and c), this overall behavior persists, although the eigenvalues always remain distinct, and there is no true bifurcation. Thus in part A, as crosstalk increases, the eigenvalues at first approach each other, and the solution remains almost segregated. At the “pseudocritical”, value of $q = \frac{(2v+\delta)(2v+\delta-2c)-\delta^2}{2v+\delta-2c}$, the eigenvalues reach their closest value, (in an “avoided crossing”) and then start to separate as crosstalk increases further; significantly beyond this pseudocritical value, the outcome is almost unsegregated (see Figures 4 and 5). Of course, for q values very close to this pseudocritical value, desegregation is very rapidly increasing with increases in crosstalk (see Figures 4 and 5), especially with very small values of δ . Thus even with slight input bias, the overall behavior, switching from segregation to unsegregation at a critical crosstalk value, resembles that seen in the unbiased situation.

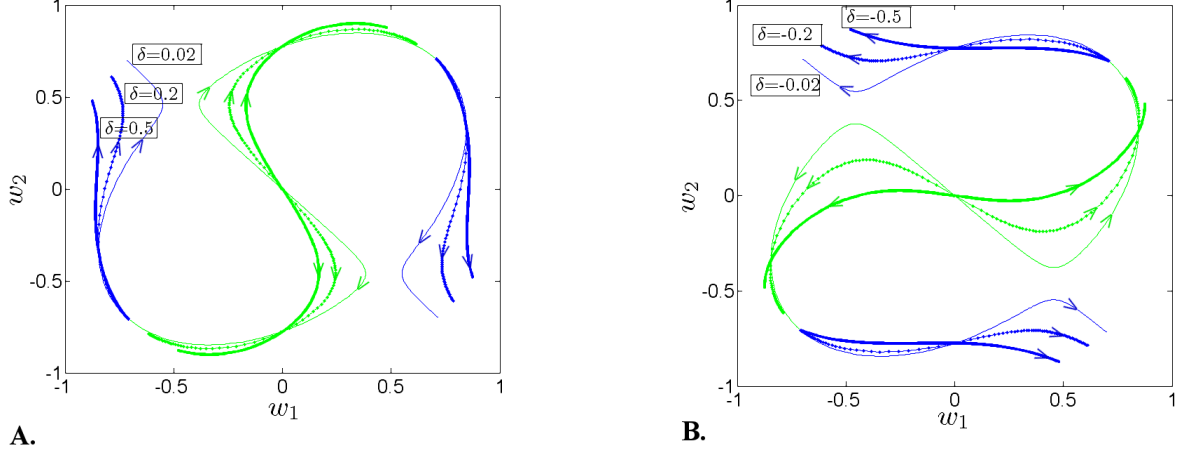


Figure 5: **Equilibria curves in the phase plane, as q changes.** The blue curves represent the stable equilibrium locus, and the green curves the saddle equilibrium. **A.** Plots for a few representative positive δ values: $\delta = 0.02$ (thin curves), $\delta = 0.2$ (thin dotted curves) and $\delta = 0.5$ (thick curves). All green saddle curves concur at one point (on the vertical axis), and all blue stable curves also concur at a point, corresponding to the fact that the position of the two equilibria is independent on the magnitude of $\delta > 0$. **B.** Plots for a few representative negative δ values: $\delta = -0.02$ (thin curves), $\delta = -0.2$ (thin dotted curves) and $\delta = -0.5$ (thick curves). All green saddle curves concur at one point, and all blue stable curves also concur at a point (on the vertical axis), corresponding to the fact that the position of the two equilibria is independent on the magnitude of $\delta < 0$. The arrows along the curves indicate the direction of increasing q .

2.4 Conclusions: mathematical behavior of the 2D system

Corollary 2.6. *For any combination of parameters, the phase-plane of the system (1) contains no cycles. Moreover, the system has only two (opposite) attracting equilibria, with attraction basins two open half-planes.*

Remark. The result holds more generally for an n -dimensional system, as shown in Appendix 1.

Since we are looking at a 2-dimensional system, this means, according to the Poincaré-Bendixon theorem, that the only attracting sets can be attracting equilibria. The two attracting equilibria of the system (by Proposition 2.3) lie along the invariant line corresponding to the largest eigenvalue of the covariance matrix \mathbf{C} , hence their position (direction and distance to origin) depend on the values of the parameters (in particular on the quality q and bias factor δ). Figure 5 illustrates the evolution of these points in the phase plane for a fixed $\delta \neq 0$, as q increases. (We used Matcont continuation algorithms to numerically estimate the equilibria and draw the equilibrium curves.)

The following two paragraphs summarize the conclusions obtained throughout the previous sections:

Biased dynamics. *When the system is biased (i.e., $\delta \neq 0$) the two eigenvalues of the input covariance matrix \mathbf{C} are always separated. The phase plane has two pairs of nonzero opposing equilibria, each situated on one of*

two distinct invariant lines through the origin (i.e., the two eigendirections of \mathbf{EC}). The invariant line of slope z_1 corresponding to the higher eigenvalue μ_1 of \mathbf{EC} contains the pair of opposing attracting equilibria; the invariant line of slope z_2 corresponding to the lower eigenvalue μ_2 of \mathbf{EC} separates their two basins of attraction and also contains the pair of opposing saddles. As the parameter q increases, the invariant lines rotate (clockwise if $\delta > 0$ and counter-clockwise if $\delta < 0$) in a continuously differentiable manner, with an angular speed that depends on q . This rotation gets arbitrarily fast (e.g., at its point of maximal rotational speed) as $\delta \rightarrow 0$.

Unbiased dynamics. When the system is unbiased (i.e., $\delta = 0$) the two eigenvalues of the input covariance matrix \mathbf{C} collide at the critical value of the quality parameter $q = q^*$. For any $q \neq q^*$, the phase plane has two pairs of nonzero opposing equilibria, each situated on one of two distinct invariant lines $\begin{pmatrix} 1 \\ \pm 1 \end{pmatrix}$ through the origin. The invariant line corresponding to the higher eigenvalue of \mathbf{EC} contains the pair of opposing attracting equilibria; the invariant line corresponding to the lower eigenvalue of \mathbf{EC} separates their two basins of attraction and also contains the pair of opposing saddles. As the parameter q increases, the invariant lines remain unchanged, until they swap instantaneously as q traverses the critical state $q = q^*$ (stability-swapping bifurcation). At the bifurcation point, the phase plane has an entire ellipse of half-stable equilibria.

Remark. The **codimension 2 bifurcation** that occurs at $q = q^*$ in the slice $\delta = 0$ can be considered a limit case of the phase-plane transition sequence obtained when increasing q , when making $\delta \rightarrow 0$ in the biased case. The rotational speed blows up to ∞ as $\delta \rightarrow 0$, and, in the $\delta = 0$ slice, the rotation becomes instantaneous via what appears to be the bifurcation’s “swap” of eigendirections. The evolution of the rotation speed with respect to q as $\delta \rightarrow 0$ is further illustrated in Figure 6.

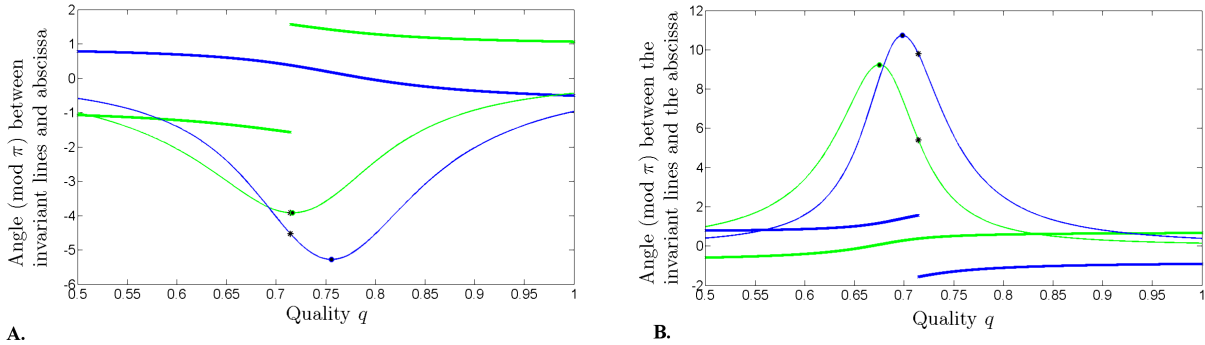


Figure 6: **Illustration of the evolution of the angles $\theta_{1,2}$ of the invariant lines with the abscissa, as q increases.** In both panels, $v = 1$ and $c = -0.4$. **A.** $\delta = 0.5$; **B.** $\delta = -0.2$. The graphs of the functions are shown in thick lines, θ_1 in blue and θ_2 in green. The graphs of the derivatives are plotted in thin lines, with $d\theta_1/dq$ in blue and $d\theta_2/dq$ in green. On the graphs of the derivatives, we marked with a black star the points corresponding to $q = q^*$, and with a bullet the points of extremum (the inflection points for $\theta_{1,2}$, where the rotational speed is maximal).

3 Alternative models: Euclidean normalization of weights versus the Oja model

The Oja rule is an elegant and classical solution to the well-known problem that unconstrained Hebbian learning is unstable [13, 33]. It has the biologically appealing feature that it is local, although it does require, somewhat implausibly, that the “normalizing” adjustment is proportional to the current weight. We have shown that it is still useful when some crosstalk is present, although the stable norm, and the exact direction of the learned weight vector, changes. One can imagine various other ways, possibly involving “homeostasis” or “synaptic scaling” [43, 44] of promoting stability, and some studies invoke various combinations of these mechanisms. A less biologically plausible, nonlocal but extremely simple and highly effective method, which might capture features of any more plausible scheme and which works even for nonlinear rules, is to impose a specific norm after each weight vector update. Here

we examine how crosstalk affects such “explicit” or “brute” normalization.

As before, Hebb’s rule lies at the basis of the weight updates: $\Delta \mathbf{w} = \gamma y \mathbf{x}$, with $y = \mathbf{w}^t \mathbf{x} = \mathbf{x}^t \mathbf{w}$.

In other words: $\mathbf{w}(n+1) = \mathbf{w} + \gamma y \mathbf{x}$. As in the Oja model, we can think of Hebbian inspecificity being formalized as a stochastic error matrix \mathcal{E} , so that, at each time step:

$$\mathbf{w} \rightarrow \mathbf{w} + \gamma y \mathcal{E} \mathbf{x}$$

Taking expectation of both sides and re-naming $\mathbf{w} = \langle \mathbf{w} \rangle$ (the long-term average of the weight vector), $\mathbf{C} = \langle \mathbf{x}^t \mathbf{x} \rangle$ (the correlation matrix of the input distribution) and $\mathbf{E} = \langle \mathcal{E} \rangle$ (the average error matrix), we obtain the iteration: $\mathbf{w} \rightarrow \mathbf{w} + \gamma \langle \mathcal{E} \mathbf{x} \mathbf{x}^t \rangle \mathbf{w} = \mathbf{w} + \gamma \mathbf{E} \mathbf{C} \mathbf{w}$.

We normalize to keep $\|\mathbf{w}\| = 1$, and make no further approximations to implement this normalization biologically. We get the new iteration function that describes the average iterative process, with errors, becomes:

$$f(\mathbf{w}) = \frac{\mathbf{w} + \gamma \mathbf{E} \mathbf{C} \mathbf{w}}{\|\mathbf{w} + \gamma \mathbf{E} \mathbf{C} \mathbf{w}\|}$$

where the “modified” covariance matrix is as before $\mathbf{E} \mathbf{C}$; unlike in the Oja case, $\mathbf{E} \mathbf{C}$ is now involved in the normalization step as well. Notice that, since $\mathbf{E} \mathbf{C}$ has positive eigenvalues, the matrix $\mathbf{I} + \gamma \mathbf{E} \mathbf{C}$ is nonsingular, hence f is defined for all $\mathbf{w} \in \mathbb{R}^n \setminus \{0\}$. The rest of the section is dedicated to discussing the position and stability of the equilibria of this new system, in whose case the direct normalization confines the trajectories to the unit circle.

In order to slightly simplify the notation, we call $\mathbf{A} = \mathbf{E} \mathbf{C}$, $\mathbf{u} = \mathbf{w} + \gamma \mathbf{E} \mathbf{C} \mathbf{w}$ and $a = \|\mathbf{u}\|$, notation which we will use whenever it is convenient. We want to see if the long-term evolution of \mathbf{w} predicted by this model is comparable with the behavior of our stochastic, discrete simulations for case where the “ratio” normalization was replaced by a “subtractive” Taylor approximation of it (see also [36]).

The vector \mathbf{w} is a fixed point of $f(\mathbf{w})$ iff $\mathbf{w} + \gamma \mathbf{A} \mathbf{w} = a \mathbf{w}$, i.e. \mathbf{w} is a unit eigenvector of \mathbf{A} (with the Euclidean norm). To establish the stability, we compute the Jacobian matrix of f at each fixed point.

Fix $j \in \overline{1, n}$. Then, for any $i \neq j$:

$$\frac{\partial u_i}{\partial w_j} = \frac{\partial}{\partial w_j} (w_i + \gamma [\mathbf{A} \mathbf{w}]_i) = \gamma A_{ij}$$

When $i = j$, we have similarly:

$$\frac{\partial u_j}{\partial w_j} = \frac{\partial}{\partial w_j} (w_j + \gamma [\mathbf{A} \mathbf{w}]_j) = 1 + \gamma A_{jj}$$

Hence, overall:

$$\frac{\partial}{\partial w_j} \|\mathbf{u}\|^2 = 2u_j(1 + \gamma A_{jj}) + \sum_{i \neq j} 2u_i \gamma A_{ij} = 2u_j + 2\gamma \sum_i u_i A_{ij} = 2u_j + 2\gamma [\mathbf{A}^t \mathbf{u}]_j$$

In matrix form:

$$\frac{\partial}{\partial \mathbf{w}} \|\mathbf{u}\|^2 = 2\gamma \mathbf{A}^t \mathbf{u} + 2\mathbf{u} \quad (4)$$

Now, fix $i \in \overline{1, n}$. For $j \neq i$, we have:

$$\frac{\partial f_i}{\partial w_j} = \frac{\gamma A_{ij} \|\mathbf{u}\| - u_i \|\mathbf{u}\|^{-1} [\gamma \mathbf{A}^t \mathbf{u} + \mathbf{u}]_j}{\|\mathbf{u}\|^2}$$

For $j = i$, we have:

$$\frac{\partial f_i}{\partial w_i} = \frac{(1 + \gamma A_{ii})\|\mathbf{u}\| - u_i\|\mathbf{u}\|^{-1}[\gamma \mathbf{A}^t \mathbf{u} + \mathbf{u}]_i}{\|\mathbf{u}\|^2}$$

Rewritten in matrix form:

$$\frac{\partial f}{\partial \mathbf{w}} = \frac{\gamma}{a} \mathbf{A} - \frac{1}{a^3} (\gamma \mathbf{u} \mathbf{u}^t \mathbf{A} + \mathbf{u} \mathbf{u}^t) + \frac{1}{a} \mathbf{I} = \frac{1}{a} \left(\mathbf{I} - \frac{1}{a^2} \mathbf{u} \mathbf{u}^t \right) (\gamma \mathbf{A} + \mathbf{I}) \quad (5)$$

where \mathbf{I} is the appropriate size identity matrix.

At any fixed point \mathbf{w} , for which automatically $\|\mathbf{w}\| = 1$ and $\mathbf{A}\mathbf{w} = \lambda_{\mathbf{w}}\mathbf{w}$, where $1 + \lambda_{\mathbf{w}}\gamma = a$), we have that:

$$\mathbf{u} \mathbf{u}^t = (\mathbf{w} + \gamma \mathbf{A}\mathbf{w})(\mathbf{w} + \gamma \mathbf{A}\mathbf{w})^t = (1 + 2\gamma\lambda_{\mathbf{w}} + \gamma^2\lambda_{\mathbf{w}}^2)\mathbf{w}\mathbf{w}^t = (1 + \lambda_{\mathbf{w}}\gamma)^2\mathbf{w}\mathbf{w}^t$$

The Jacobian at a fixed point \mathbf{w} can be then simplified to:

$$\frac{\partial f}{\partial \mathbf{w}} = \frac{1}{a} (\mathbf{I} - \mathbf{w}\mathbf{w}^t) (\gamma \mathbf{A} + \mathbf{I}) \quad (6)$$

We calculate:

$$\frac{\partial f}{\partial \mathbf{w}}(\mathbf{w}) = \frac{1}{a} (\mathbf{I} - \mathbf{w}\mathbf{w}^t) (\gamma \lambda_{\mathbf{w}} + 1) \mathbf{w} = \frac{1}{a} (\mathbf{w} - \mathbf{w}(\mathbf{w}^t \mathbf{w})) (\gamma \lambda_{\mathbf{w}} + 1) = 0$$

Complete \mathbf{w} to a basis of eigenvectors of \mathbf{A} (not necessarily mutually orthogonal). Let $\mathbf{v} \neq \mathbf{w}$ any of the vectors in this basis (with eigenvalue $\lambda_{\mathbf{v}}$), and consider $\mathbf{z} = \mathbf{v} - [\mathbf{w}^t \mathbf{v}] \mathbf{w}$ the projection of \mathbf{v} on the orthogonal complement of \mathbf{w} . Then:

$$(\gamma \mathbf{A} + \mathbf{I})(\mathbf{z}) = (\gamma \lambda_{\mathbf{v}} + 1) \mathbf{v} - (\gamma \lambda_{\mathbf{w}} + 1) [\mathbf{w}^t \mathbf{v}] \mathbf{w}$$

Hence

$$\frac{\partial f}{\partial \mathbf{w}}(\mathbf{z}) = \frac{1}{a} (\gamma \lambda_{\mathbf{v}} + 1) (\mathbf{v} - [\mathbf{w}^t \mathbf{v}] \mathbf{w}) = \frac{1}{a} (\gamma \lambda_{\mathbf{v}} + 1) \mathbf{z} = \frac{\gamma \lambda_{\mathbf{v}} + 1}{\gamma \lambda_{\mathbf{w}} + 1} \mathbf{z}$$

A normalized eigenvector \mathbf{w} of \mathbf{EC} is stable as a fixed point of the system if all the eigenvalues of the Jacobian $\frac{\partial f}{\partial \mathbf{w}}$ at \mathbf{w} are less than one in absolute value :

$$\left| \frac{\gamma \lambda_{\mathbf{v}} + 1}{\gamma \lambda_{\mathbf{w}} + 1} \right| < 1$$

Since all eigenvalues of \mathbf{EC} are positive (recall that \mathbf{EC} is diagonalizable with the dot product $\langle \cdot, \cdot \rangle_{\mathbf{C}}$), this is equivalent to $\frac{\gamma \lambda_{\mathbf{v}} + 1}{\gamma \lambda_{\mathbf{w}} + 1} < 1$, and thus to $\lambda_{\mathbf{w}} > \lambda_{\mathbf{v}}$ for every $\mathbf{v} \neq \mathbf{w}$.

In conclusion: The system has stable fixed points iff the modified correlation matrix \mathbf{EC} has a maximal eigenvalue of multiplicity one. Then, a point \mathbf{w} is a stable fixed point of the system iff it is a unit eigenvector of \mathbf{EC} corresponding to the unique maximal eigenvalue of \mathbf{EC} .

It is now clear that the phase space of this system, although not dynamically equivalent to the phase space of the corresponding Oja model, it is very similar. Disregarding the origin (which is not in the domain of one, but is a repelling fixed point for the other), the other fixed points have the same qualitative behavior (stability) for both systems, if assuming γ sufficiently small. Moreover, the stability transitions occur at the same bifurcation points (where the eigenvalues of $b f E C$ collide with each other), and the bifurcation phase-planes are themselves similar.

In the case of the two-dimensional model discussed in this paper, the eigenvalue swap occurs as before when $q = q^* = \frac{v}{v-c}$. For example, in the unbiased case $\delta = 0$, the bifurcation phase plane at $q = q^*$ again exhibits an ellipse attractor. Indeed, at the codimension 2 $q = q^*$, the iteration function becomes: $f(\mathbf{w}) = \frac{\mathbf{w}}{\|\mathbf{w}\|}$, which maps radially any \mathbf{w} in the plane to the unit circle, and keeps it fixed thereafter.

The next section shows phase plane simulations for both models, in the more realistic situation of stochastic weight updates, in discrete time and at finite learning rate.

4 Stochastic models. Simulations and predictions

Here we briefly study the more biologically realistic situation in which the weights update stochastically with a small finite learning rate γ , driven by each individual input, rather than by the mean statistics in the negligible learning rate limit. In particular we study whether the convergence to eigenvector equilibria, and the transitions in dynamics between different values of the parameter q still occur as in the deterministic model.




Figure 7: **Behavior of stochastic weight updates for a biased input distribution $\delta \neq 0$.** **A.** A discrete input sample ($N = 4000$) was drawn out of an input distribution with $v = 1$, $c = -0.4$, $\delta = 1$, and used to update the weights. **B.** Depending on their initial state, the weight vector stabilizes towards small stochastic fluctuations around either one of the attracting equilibria (the pair of appropriately normalized eigenvectors corresponding to the larger eigenvalue of \mathbf{EC}). **C.** The corresponding iterations are shown in the case of exact normalization at each step (fewer iterations are shown in this case, since more weights, all living on the unit circle, would obstruct the clarity of the figure.) In all three panels, the points were colored update-chronologically from red to blue. We used the critical quality $q = 1/1.4 \sim 0.71$.

Our numerical simulations show, as expected, that convergence is conserved, in the following sense: when a pair of attracting equilibria exist for the deterministic system (i.e., \mathbf{EC} has distinct eigenvalues), the discrete sequence of updating $\boldsymbol{\omega}$ eventually stabilizes to small, stochastic fluctuations around one of these two equilibria (which are, as we recall, the appropriately normalized eigenvectors corresponding to the larger eigenvalue of \mathbf{EC}). This is illustrated in Figures 8, 9a and 9c. In Figure 7, \mathbf{x} is drawn out of a biased distribution of inputs (shown on the left), for which the two eigenvalues of \mathbf{EC} are warranted to be distinct for any value of q , in particular for the value chosen here ($q = q^* = 1/1.4$). In Figure 8, the inputs are unbiased, so the same remark applies only if $q \neq q^*$. In Figure 7a, we illustrate the case $q > q^*$, in which the attracting vectors are $\pm \sqrt{q-1/2} \begin{pmatrix} 1 \\ -1 \end{pmatrix}$; in Figure 7c, we illustrate the case $q < q^*$, in which the attracting vectors are $\pm \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$. In both cases, the stochastic update settles to fluctuations about either one of these vectors, depending on the initial conditions.

Figure 8b illustrates the unbiased case corresponding to the codimension 2 bifurcation in the deterministic dynamics; that is, when $q = q^*$. Recall that, in the deterministic phase-plane, this case was characterized by an ellipse of neutrally stable equilibria, so that each initial condition would converge radially towards a unique nonisolated equilibrium on this curve. This situation changes in the model driven by stochastic updates. An initial weight vector $\boldsymbol{\omega}$ will quickly be attracted towards the ellipse; however, the orbit does not fluctuate around a particular point on the curve, but rather perpetually drifts along the curve, eventually covering densely the entire ellipse.

5 Discussion

We have proposed [11, 1] that a central problem for biological learning is that the activity-dependent processes that lead to connection strength adjustments cannot be completely synapse specific, because they must obey the laws of physics. This truism provides a new viewpoint: it raises the possibility that sophisticated learning, such as presumably occurs in the neocortex, is enabled as much by special machinery for enhancing specificity, as by special algorithms [4]. We have suggested that these plasticity errors are analogous to mutations, and that cortical circuitry might reduce such errors, just as “proofreading” reduces dna copying mistakes. In particular, it seems possible that the key to overcoming the curse of dimensionality that underlies difficult, and apparently almost intractable, learning problems lies not just in finding good approximations, architectures and techniques, but also in perfecting

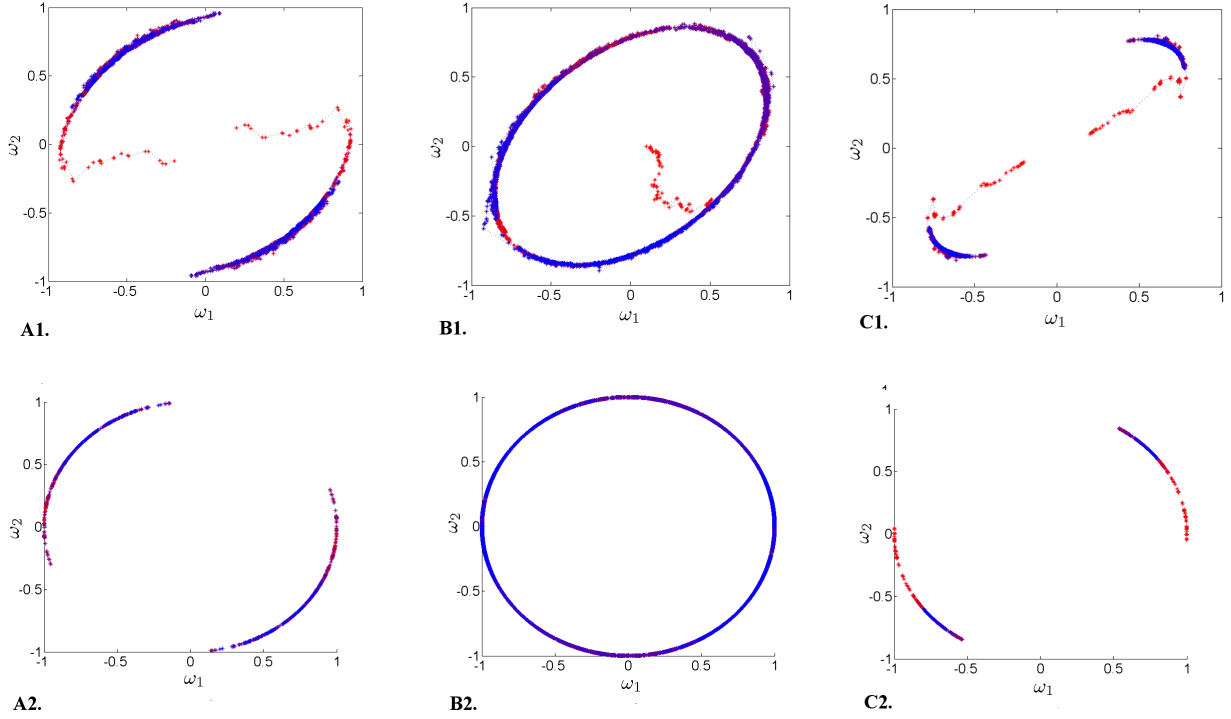


Figure 8: **Differences in stochastic behavior when q is varied, in the unbiased input case $\delta = 0$ (compare with Figure 3).** A discrete number of input vectors $\mathbf{x}(t) = (x_1(t), x_2(t))$ are drawn from a distribution with covariance matrix C , with $v = 1$, $c = -0.4$ (so that the critical quality value $q^* = 1/1.4 \sim 0.71$). The weights $\omega = (\omega_1(t), \omega_2(t))$, adjusting with a small learning rate $\gamma = 0.1$, are plotted in the (ω_1, ω_2) plane, with the color of the points changing chronologically from red to blue. The top panels show the behavior of the Oja model, while the bottom panels, for corresponding parameters, show the behavior for exact normalization of weights at each step. **A.** For good transmission quality $q = 0.85 > q^*$, ω is converging in the long term to a state of small fluctuations around either $\pm \sqrt{q - 1/2} (1 \ -1)^T$, depending on the initial state. The plot illustrates the trajectories for two initial states, each stabilizing around one of these opposite eigenvectors. **B.** For critical transmission quality $q = q^*$, ω converges to fluctuations around the ellipse of neutrally attracting equilibria, but will perpetually drift around, filling the ellipse, driven by input fluctuations from the mean statistics, without remaining asymptotically near any particular equilibrium state. **C.** For poor transmission quality $q = 0.6 < q^*$, ω is converging in the long term to a state of small fluctuations around $\pm \frac{1}{\sqrt{2}} (1 \ 1)^T$, depending on the initial state. The plot illustrates the trajectories for two initial states, each stabilizing to stabilizing around one of these opposite eigenvectors.

the relevant biological plasticity apparatus. Indeed, it seems possible that problems of survival and reproduction are so diverse that no single algorithm can solve them all, so that no “universal” or “canonical” cortical circuit would be expected. In these circumstances, as Rutherford once said about physics, neuroscience would become a type of stamp collecting. However, if every specialized algorithm relies on extraordinarily specific synaptic weight adjustment, then finding machinery that allows such specificity would indeed be tantamount to discovering new neurobiological general principles, somewhat along the lines that established the main framework for modern biology (Darwinian evolution, Mendelian genetics, DNA structure and function, replication mechanisms etc). We have speculated that an important part of such machinery, at least in the neocortex, might lie outside the synapse itself, in the form of complex circuitry performing a proofreading operation analogous to that procuring accuracy for polynucleotide copying [1, 4, 2]. However, such machinery would be less necessary if update inaccuracy merely degraded learning, rather than preventing it. In particular even if temporarily unfavorable (e.g., “noisy”) input statistics led to imperfect learning because of Hebbian inspecificity, the degraded weights might still be a useful starting point for better learning when input statistics improve. On the other hand, if inspecificity completely prevented even partial learning, then

rapid and successful learning from newly favorable statistics might be impossible. These considerations have impelled us to examine the effect of Hebbian “crosstalk” in various classical models of unsupervised learning, using both linear [36] and nonlinear rules [11] (see also [17]).

5.1 Separate but equal: segregation without bias

In this paper we extended our previous study [36] of the effect of crosstalk on the simple linear Hebbian model of Oja to situations approaching the “unbiased” case where all inputs have the same statistical distribution. This case has often been invoked in discussions of the emergence of ocular dominance wiring and other forms of neural development, but it might also apply to any situation in which sets of inputs disconnect completely, or “segregate,” to form pruned wiring patterns that are then “sculpted” by a more subtle synaptic learning process (of course in the present model weights and activities can be negative; we interpret negative weights as disconnections). For the case of visual input, it seems likely that statistics would be similar, and positively correlated, for the two eyes, which look at the same world, and it is well known [13] that a linear Hebb rule with unbiased inputs, under either implicit or explicit normalization, leads to the symmetric, equal-weight, and thus apparently unbiological, outcome. A possible but rather unbiological solution to this is to use a “subtractive” normalization scheme, although this also requires imposing weight limits [32]. It has been shown that a wide variety of nonlinear rules [15], including the BCM rule [8] and STDP [16] can lead to ocular segregation under unbiased statistics. The key point is that segregated states can be created by typically nonlinear, “symmetry-breaking” mechanisms even when the inputs themselves do not favor particular segregated outcomes. Indeed, the absence of bias could be characteristic of development, as opposed to “learning,” insofar as these two notions are distinct.

A natural question would be: if such segregated outcomes are an important part of normal development (which then constrain subsequent, more detailed, “refining,” plasticity processes, including learning), how could the determining “unbiased” statistics arise, and conversely, how would plasticity errors, such as crosstalk, or other alterations in the form of the rule, affect the outcome? In particular we show here that, unsurprisingly, crosstalk tends to prevent segregation, especially when the inputs are close to unbiased. This might set a limit to the use of symmetry breaking to generate specific wiring, or require special specificity-enhancing circuitry, such as “proofreading”, even during development. At the very least it suggests that internally generated patterns deriving segregation, such as negative correlations, might have to be quite strong to overcome the desegregating effect of inevitable crosstalk.

Before exploring this further, we comment briefly about “unbias” in relation to Hebbian learning. Although here we focus on lack of bias in the second order statistics, one can also postulate unbias at all order, an assumption which greatly simplifies the study of nonlinear Hebbian plasticity, essentially eliminating the possibility of learning and restricting analysis to development. To what would unbiased high-order statistics correspond? It seems that they correspond to the radially symmetric distributions recently considered by Lyu and Simoncelli [28], where the joint pdf equal density contour lines are nested hyperspheres with nonGaussian spacings. One might expect that with completely unbiased (spherical) input statistics no particular direction in weight space would be favored and therefore the outcomes would be either symmetric (equal weights), or broken symmetric (various combinations of opposite but equal magnitude weights); the particular set of outcomes would be determined by the higher-order correlations, and could be quite complicated. Indeed, Elliott [15] finds that segregated outcomes are quite typical of nonlinear Hebbian rules with unbiased statistics and shows that crosstalk can induce bifurcations in these cases [17].

Recently, it has been suggested that the Oja rule (even without crosstalk) and Eigen’s replication/mutation equation might be “isomorphic” [20, 19]. Indeed both equations describe normalized growth processes. However, our work shows that the Oja equation only shows a bifurcation at a critical crosstalk value in very narrow conditions. We suggest that the important analogy lies less in detailed mathematical equivalencies, and more in the fundamental need for accuracy in elementary biological processes. In particular, it’s clear that superaccurate polynucleotide copying underlies Darwinian evolution, and similarly superaccurate Hebbian plasticity might be needed for neural learning.

5.2 Effect of crosstalk on linear learning

The analysis reported here essentially shows that the well known bifurcation that occurs in linear Hebbian learning as unbiased negative correlations become positive (from segregated to unsegregated states) still occurs in the presence of crosstalk, but at a new, crosstalk-dependent critical negative correlation level. This effect is quite intuitive: crosstalk favors the unsegregated state, and therefore allows the switch to occur at negative correlation, rather than at zero correlation. Of course this situation changes dramatically as soon as any degree of bias is introduced, since now the eigenvalues of \mathbf{C} become distinct, and our previous analysis [36] applies : crosstalk produces a smooth change in the

direction of the learned weights (the dominant eigenvector of \mathbf{EC}). Our present analysis attempts to characterize the relation between these two regimes. In particular, we show that the smooth change can be very rapid when bias is weak.

The change in the normalization produced by crosstalk in the Oja model is largely irrelevant, and indeed one still sees the same behavior with explicit normalization (Section 3). Our analysis also gives insight into the codimension two bifurcation that occurs at the critical quality q^* , via an ellipse of half-stable equilibria. The motion towards the ellipse becomes extremely rapid (Figures 3B and 9), which permits the exchange of stability between the 2 invariant lines (Figure 2). This rapid motion shows up in simulations (and presumably in biological realizations) as very “noisy” weights as the threshold crosstalk value is neared.

Of course a true bifurcation is only seen for unbiased inputs and for negligible learning rates. However, the behavior remains practically indistinguishable from a bifurcation even with slightly biased inputs and finite learning rates. An example was already discussed in our previous paper (see Figure in [36]). A similar situation occurs with models of phase transitions: a true bifurcation of the dynamics only occurs in the “thermodynamic limit,” but this is effectively established even for quite small systems [39]. We have previously called attention to the analogy between Hebbian learning and molecular evolution [3, 19], with crosstalk playing the role of mutation. In Eigen’s evolution model [14], the transition from the ordered, living, state to the disordered, chemical, state is quite sharp even for polynucleotide lengths ~ 50 , though a true phase transition (identical to that of the surface of the 2-dimensional Ising model) is only seen with unlimited chains [37]. Interestingly, the model becomes easiest to analyze in this limit, and the relevant dimensionless control parameter is simply the product of the mutation rate and the (binary) chain length (for binary strings). Although we analyzed here the $n = 2$ case, we assume that weights are specified with unlimited bit resolution (i.e., reals). In this case the dimensionless control parameter, equivalent to that in the thermodynamic limit of the Eigen model, is q . In the standard Eigen model, the mutation rate is the same at all chain positions. In the next section we discuss the analogous concept for Hebbian learning.

5.3 The error matrix

Throughout this paper we assume that the Hebbian adjustment of any weight was equally affected by error, and does not depend either on the strength of that weight or its identity. Such “isotropicity” seems a reasonable first assumption, like neglecting bumps on an inclined plane in mechanics. However, it does appear to fly in the face of biological reality. First, stronger synapses are also bigger, and also require higher spine neck conductances and therefore presumably are less well isolated chemically [26]. However, such expected “weight-dependence” might only be a second-order effect, because in order to ensure that LTP is “Hebbian,” the spine neck resistance must always be sufficiently low, even at “silent,” AMPAR-less zero-strength synapses, so that the essential back-propagating spike effectively invades the spine head [29] during the peak NMDAR opening. Second, crosstalk between individual synapses is a relatively local, not global, phenomenon [9, 23]. However, during learning individual synapses appear and disappear, which will smear details. Furthermore connections are made up of many individual synapses scattered over much of the dendritic tree, which will also smear detail [36]. Recent work shows that feed-forward cortical connections carrying similar information do not “cluster” on dendritic segments, invalidating the argument that local crosstalk could promote useful clustering [23].

5.4 Relevance to ocular dominance and general developmental mechanisms

A useful though rather fuzzy distinction can be drawn between developmental mechanisms which generate sets of connections (“circuits”), or, perhaps, “incipient” or “potential” connections [1, 40] that can be made actual without axo-dendritic rewiring merely by adding postsynaptic spines or presynaptic “drumsticks” [5, 38], and “learning,” which refines (perhaps in crucial ways) the overall framework established by development. This distinction is related to that between “Nature” and “Nurture,” or, in the context of Chomskyan linguistics, “principles” and “parameters.” The Oja model encapsulates this distinction in minimal form: by definition when the inputs are unbiased there can be no learning, and only two outcomes are possible, which we call segregated or unsegregated. The classic biological example is that in many species early in development a geniculate axon diffusely innervates a patch of layer IV of cortex (though it does not necessarily contact all the neurons whose dendrites ramify in that patch), but then retracts from stripes within that patch that become selectively innervated by axons corresponding carrying signals from the other eye. Cells within a stripe then becomes largely monocular, although they develop different selectivities for different stimulus features such as orientation. In the Oja model segregation appears in response to unbiased (or, effectively, nearly unbiased) inputs at a critical level of negative correlation, which depends on the degree of

crosstalk. In real animals segregation appears before the onset of visual experience, and is thought to be driven by unbiased inputs generated by spontaneous firing. While one might expect crosstalk to hinder segregation (since it tends to equalize weights), our results show this is not quite correct in the Oja model: it merely shifts the critical degree of (unbiased) correlation required. Various proposals exist for how such inputs can induce segregation even when correlations are positive [32, 15] and it’s likely that crosstalk will also have the same weak effect here. Indeed, Elliott [17] has shown that while crosstalk induces a bifurcation from segregation to unsegregation in a weight-dependent model, the critical value (his equation 3.8) can be shifted to favorable values with suitable correlation values. Thus, the endogenous developmental machinery that creates circuits probably does not require great Hebbian accuracy (and might not require Hebbian machinery at all [12, 35]). If the aforementioned postulated layer VI proofreading circuit [4] underlies accuracy, it would not be needed until learning begins, consistent with evidence that the final stages of layer VI circuitry (for example, feedback to relay cells) is late to develop. Indeed, much of the initial pruning that takes place in development might serve to improve the accuracy of proofreading circuitry essential for true learning.

Once detailed, and biased, sensory input occurs, it can drive quantitative adjustments in the already correctly segregated circuits, involving both synapse-strength change and stabilization and un-silencing of new spines (and removal of weak synapses). However, even in the highly simplified Oja model, appropriate adjustment now requires great accuracy, and therefore presumably “proofreading,” especially when correlation bias is weak.

5.5 Normalization and error

We have assumed in both our papers on the linear Hebb rule that the effect of crosstalk is solely on the Hebbian part of the rule, not the normalizing component. However, if normalization, or some other process that stabilizes Hebbian learning, is biologically necessary, then it is presumably also subject to imperfections such as crosstalk. There are basically two ways to add this other form of crosstalk to stabilized Hebbian learning rules. In the context of a single neuron model, one could simply apply a second, different, crosstalk matrix, say \mathbf{F} , to the stabilizing term. However, if $\mathbf{F} = \mathbf{E}$ (because the geometry underlying such errors is the same in each case) such normalization crosstalk cancels the overall effect. Even if the pattern of crosstalk at each update is described by fluctuating matrices \mathcal{E} and \mathcal{F} , whose averages equal \mathbf{E} and \mathbf{F} , which do not exactly cancel, one might expect they would on average. Could this be a way to eliminate errors?

The other possible way that errors could affect normalization would be if the underlying normalization mechanism were sufficiently different that the average geometry differed. Although the Oja model allows negative activities, firing rates can only be positive, and it is tempting to suppose that negative signals are carried in special “off” channels whose positive weights represent negative ones. In this equivalence, the Hebb part of the rule reflect LTP and the normalization part, LTD. In the cortex, LTP seems to be postsynaptic and LTD presynaptic. If the update leaks presynaptically, within the axon, it will affect a different set of synapses, ones that mostly form onto a different postsynaptic cell. To evaluate how such errors might affect learning, one needs a multi-unit model.

6 Conclusion

Generically the inspecific Oja rule does not show bifurcations with variation in the crosstalk parameter. In this paper we analyze an interesting special case which does show a bifurcation: when the input statistics are unbiased. We also describe the behavior in the vicinity of this special case, which is practically indistinguishable from a bifurcation. Essentially in this region “learning” changes rather abruptly from being dominated by second-order input statistics (at sufficiently low crosstalk) to being dominated by the internal pattern of crosstalk itself. However, we regard this behavior as being biologically rather uninteresting, since synaptic mechanisms are presumably accurate enough that it never occurs. The one exception would be during development, where near-unbiased statistics might be used by the brain to induce initial selective wiring. Our results suggest that even in this case, high Hebbian accuracy might be required. However, extreme accuracy is probably most essential for nonlinear learning from higher-order statistics [11, 17].

References

- [1] P. Adams and K. Cox. A new interpretation of thalamocortical circuitry. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1428):1767, 2002.

- [2] P. Adams and K. Cox. From life to mind: two prozaic miracles. In *Integral Biomathics*, volume 67, page 02. 2012.
- [3] P.R. Adams and K.J.A. Cox. Synaptic darwinism and neocortical function. *Neurocomputing*, 42(1-4):197–214, 2002.
- [4] PR Adams and KJA Cox. A neurobiological perspective on building intelligent devices. *Neuromorphic Eng*, 3:2–8, 2006.
- [5] J.C. Anderson and K.A.C. Martin. Does bouton morphology optimize axon length? *nature neuroscience*, 4(12):1166–1167, 2001.
- [6] A.J. Bell and T.J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159, 1995.
- [7] G.Q. Bi. Spatiotemporal specificity of synaptic plasticity: cellular rules and mechanisms. *Biological Cybernetics*, 87(5):319–332, 2002.
- [8] E.L. Bienenstock, L.N. Cooper, and P.W. Munro. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *The Journal of Neuroscience*, 2(1):32–48, 1982.
- [9] T. Bonhoeffer, V. Staiger, and A. Aertsen. Synaptic plasticity in rat hippocampal slice cultures: local “hebbian” conjunction of pre-and postsynaptic stimulation leads to distributed synaptic enhancement. *Proceedings of the National Academy of Sciences*, 86(20):8113, 1989.
- [10] L.N. Cooper. *Theory of cortical plasticity*. World Scientific Pub Co Inc, 2004.
- [11] K.J.A. Cox and P.R. Adams. Hebbian crosstalk prevents nonlinear unsupervised learning. *Frontiers in computational neuroscience*, 3, 2009.
- [12] J.C. Crowley and L.C. Katz. Early development of ocular dominance columns. *Science*, 290(5495):1321, 2000.
- [13] P. Dayan and LF Abbott. Theoretical neuroscience: Computational and mathematical modeling of neural systems. 2001.
- [14] M. Eigen. Selforganization of matter and the evolution of biological macromolecules. *Naturwissenschaften*, 58(10):465–523, 1971.
- [15] T. Elliott. An analysis of synaptic normalization in a general class of hebbian models. *Neural computation*, 15(4):937–963, 2003.
- [16] T. Elliott. Temporal dynamics of rate-based synaptic plasticity rules in a stochastic model of spike-timing-dependent plasticity. *Neural computation*, 20(9):2253–2307, 2008.
- [17] T. Elliott. Cross-talk induces bifurcations in nonlinear models of synaptic plasticity. *Neural Computation*, pages 1–68, 2012.
- [18] T. Elliott and N.R. Shadbolt. Multiplicative synaptic normalization and a nonlinear hebb rule underlie a neurotrophic model of competitive synaptic plasticity. *Neural computation*, 14(6):1311–1322, 2002.
- [19] C. Fernando, R. Goldstein, and E. Szathmáry. The neuronal replicator hypothesis. *Neural computation*, 22(11):2809–2857, 2010.
- [20] C. Fernando and E. Szathmáry. Chemical, neuronal and linguistic replicators. *Towards an Extended Evolutionary Synthesis Cambridge, Ma.: MIT Press*, 2009.
- [21] P. Földiák. Forming sparse representations by local anti-hebbian learning. *Biological cybernetics*, 64(2):165–170, 1990.
- [22] G.J. Goodhill. Topography and ocular dominance: a model exploring positive correlations. *Biological Cybernetics*, 69(2):109–118, 1993.

- [23] C.D. Harvey and K. Svoboda. Locally dynamic synaptic learning rules in pyramidal neuron dendrites. *Nature*, 450(7173):1195–1200, 2007.
- [24] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent component analysis*, volume 26. Wiley-Interscience, 2001.
- [25] A. Hyvärinen and E. Oja. Independent component analysis by general nonlinear hebbian-like learning rules. *Signal Processing*, 64(3):301–313, 1998.
- [26] C. Koch. *Biophysics of computation: information processing in single neurons*. Oxford University Press, USA, 2005.
- [27] R. Linsker. From basic network principles to neural architecture: Emergence of orientation columns. *Proceedings of the National Academy of Sciences*, 83(22):8779, 1986.
- [28] S. Lyu and E.P. Simoncelli. Nonlinear extraction of independent components of natural images using radial gaussianization. *Neural computation*, 21(6):1485–1519, 2009.
- [29] J.C. Magee and D. Johnston. A synaptically controlled, associative signal for hebbian plasticity in hippocampal neurons. *Science*, 275(5297):209, 1997.
- [30] C. Malsburg. Self-organization of orientation sensitive cells in the striate cortex. *Biological Cybernetics*, 14(2):85–100, 1973.
- [31] K.D. Miller, J.B. Keller, and M.P. Stryker. Ocular dominance column development: Analysis and simulation. *Science*, 245(4918):605, 1989.
- [32] K.D. Miller and D.J.C. MacKay. The role of constraints in hebbian learning. *Neural Computation*, 6(1):100–126, 1994.
- [33] E. Oja. Simplified neuron model as a principal component analyzer. *Journal of mathematical biology*, 15(3):267–273, 1982.
- [34] B.A. Olshausen et al. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.
- [35] S.B. Paik and D.L. Ringach. Retinal origin of orientation maps in visual cortex. *Nature Neuroscience*, 2011.
- [36] A. Radulescu, K. Cox, and P. Adams. Hebbian errors in learning: an analysis using the oja model. *Journal of theoretical biology*, 258(4):489–501, 2009.
- [37] D.B. Saakian and C.K. Hu. Exact solution of the eigen model with general fitness functions and degradation rates. *Proceedings of the National Academy of Sciences of the United States of America*, 103(13):4935, 2006.
- [38] S.M. Sherman and R.W. Guillery. *Exploring the thalamus*. Academic Pr, 2001.
- [39] P. Sollich. Finite-size effects in learning and generalization in linear perceptrons. *Journal of Physics A: Mathematical and General*, 27:7771, 1994.
- [40] A. Stepanyants and D.B. Chklovskii. Neurogeometry and potential synaptic connectivity. *Trends in neurosciences*, 28(7):387–394, 2005.
- [41] J. Swetina and P. Schuster. Self-replication with errors:: A model for polynucleotide replication2. *Biophysical chemistry*, 16(4):329–345, 1982.
- [42] NV Swindale. The development of topography in the visual cortex: a review of models. *Network: Computation in neural systems*, 7(2):161–247, 1996.
- [43] G.G. Turrigiano, K.R. Leslie, N.S. Desai, L.C. Rutherford, and S.B. Nelson. Activity-dependent scaling of quantal amplitude in neocortical neurons. *NATURE-LONDON*-, pages 892–895, 1998.
- [44] G.G. Turrigiano and S.B. Nelson. Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, 5(2):97–107, 2004.
- [45] D.J. Willshaw and C. Von Der Malsburg. How patterned neural connections can be set up by self-organization. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 194(1117):431–445, 1976.

Appendix 1. A few detailed proofs

Lemma. $Df_{\mathbf{w}}^{\mathbf{E}} = \gamma [\mathbf{E}\mathbf{C} - 2\mathbf{w}(\mathbf{C}\mathbf{w})^T - (\mathbf{w}^T\mathbf{C}\mathbf{w})\mathbf{I}]$

Proof. Call $g(\mathbf{w}) = (\mathbf{w}^T\mathbf{C}\mathbf{w})\mathbf{w}$, so $f^{\mathbf{E}}(\mathbf{w}) = \gamma[\mathbf{E}\mathbf{C}\mathbf{w} - g(\mathbf{w})]$

$$g_i(\mathbf{w}) = (\mathbf{w}^T\mathbf{C}\mathbf{w})w_i$$

If $i \neq j$:

$$\frac{\partial g_i}{\partial w_j}(\mathbf{w}) = \frac{\partial}{\partial w_j}(\sum_{k,l} C_{kl}w_kw_l)w_i = 2(\sum_k C_{kj}w_k)w_i = 2[\mathbf{C}\mathbf{w}]_jw_i$$

If $i = j$:

$$\begin{aligned} \frac{\partial g_i}{\partial w_i}(\mathbf{w}) &= \frac{\partial}{\partial w_i}(\sum_{k,l} C_{kl}w_kw_l)w_i + \sum_{k,l} C_{kl}w_kw_l = 2(\sum_k C_{ki}w_k)w_i + \\ &\quad + \mathbf{w}^T\mathbf{C}\mathbf{w} = 2[\mathbf{C}\mathbf{w}]_iw_i + \mathbf{w}^T\mathbf{C}\mathbf{w} \end{aligned}$$

So:

$$Dg_{\mathbf{w}} = 2\mathbf{w}(\mathbf{C}\mathbf{w})^T + (\mathbf{w}^T\mathbf{C}\mathbf{w})\mathbf{I}$$

□

Proposition 1.2. Suppose $\mathbf{E}\mathbf{C}$ has a multiplicity one largest eigenvalue. An equilibrium \mathbf{w} (i.e., by Proposition 1.1, an eigenvector of $\mathbf{E}\mathbf{C}$ with eigenvalue $\lambda_{\mathbf{w}}$, normalized so that $\|\mathbf{w}\|_{\mathbf{C}} = \lambda_{\mathbf{w}}$) is a local hyperbolic attractor for the system (1) iff it is an eigenvector corresponding to the maximal eigenvalue of $\mathbf{E}\mathbf{C}$.

Proof. Fix an eigenvector \mathbf{w} of $\mathbf{E}\mathbf{C}$, with $\mathbf{E}\mathbf{C}\mathbf{w} = \lambda_{\mathbf{w}}\mathbf{w}$. Then:

$$\begin{aligned} Df_{\mathbf{w}}^{\mathbf{E}}\mathbf{w} &= \gamma[\mathbf{E}\mathbf{C}\mathbf{w} - 2\mathbf{w}(\mathbf{C}\mathbf{w})^T\mathbf{w} - (\mathbf{w}^T\mathbf{C}\mathbf{w})\mathbf{w}] = \\ &= \gamma[-2\mathbf{w}\mathbf{w}^T\mathbf{C}\mathbf{w}] = -2\gamma\lambda_{\mathbf{w}}\mathbf{w} \end{aligned}$$

Recall that the vector \mathbf{w} can be completed to a basis \mathcal{B} of eigenvectors, orthogonal with respect to the dot product $\langle \cdot, \cdot \rangle_{\mathbf{C}}$. Let $\mathbf{v} \in \mathcal{B}$, $\mathbf{v} \neq \mathbf{w}$, be any other arbitrary vector in this basis, so that $\mathbf{E}\mathbf{C}\mathbf{v} = \lambda_{\mathbf{v}}\mathbf{v}$, and $\langle \mathbf{w}, \mathbf{v} \rangle_{\mathbf{C}} = \mathbf{w}^T\mathbf{C}\mathbf{v} = 0$. We calculate:

$$\begin{aligned} Df_{\mathbf{w}}^{\mathbf{E}}\mathbf{v} &= \gamma[\mathbf{E}\mathbf{C}\mathbf{v} - 2\mathbf{w}\mathbf{w}^T\mathbf{C}\mathbf{v} - \lambda_{\mathbf{w}}\mathbf{v}] = \\ &= \gamma[(\lambda_{\mathbf{v}} - \lambda_{\mathbf{w}})\mathbf{v} - 2\langle \mathbf{w}, \mathbf{v} \rangle_{\mathbf{C}}\mathbf{w}] = -\gamma[\lambda_{\mathbf{w}} - \lambda_{\mathbf{v}}]\mathbf{v} \end{aligned}$$

So \mathcal{B} is also a basis of eigenvectors for $Df_{\mathbf{w}}^{\mathbf{E}}$. The corresponding eigenvalues are $-2\gamma\lambda_{\mathbf{w}}$ (for the eigenvector \mathbf{w}) and $-\gamma[\lambda_{\mathbf{w}} - \lambda_{\mathbf{v}}]$ (for any other eigenvector $\mathbf{v} \in \mathcal{B}$, $\mathbf{v} \neq \mathbf{w}$). An equivalent condition for \mathbf{w} to be a hyperbolic attractor for the system (1) is that all the eigenvalues of $Df_{\mathbf{w}}^{\mathbf{E}}$ are < 0 . Since the learning rate γ and the eigenvalue $\lambda_{\mathbf{w}}$ are both > 0 , this condition is further equivalent to having $-\gamma(\lambda_{\mathbf{w}} - \lambda_{\mathbf{v}}) < 0$, for all $\mathbf{v} \in \mathcal{B}$, $\mathbf{v} \neq \mathbf{w}$. In conclusion, an equilibrium \mathbf{w} is a hyperbolic attractor if and only if $\lambda_{\mathbf{w}} > \lambda_{\mathbf{v}}$, for all $\mathbf{v} \neq \mathbf{w}$ (i.e. $\lambda_{\mathbf{w}}$ is the maximal eigenvalue, or in other words if \mathbf{w} is in the direction of the principal eigenvector of $\mathbf{E}\mathbf{C}$). □

Appendix 2. An extension to higher dimensions

Theorem. Suppose the modified covariance matrix \mathbf{EC} has a unique maximal eigenvalue λ_1 . Then the two eigenvectors $\pm \mathbf{w}_{\mathbf{EC}}$ corresponding to λ_1 , normalized such that $\|\mathbf{w}\|_{\mathbf{C}} = \lambda_1$, are the only two attractors of the system. More precisely, the phase space is divided into two basins of attraction, of $\mathbf{w}_{\mathbf{EC}}$ and $-\mathbf{w}_{\mathbf{EC}}$ respectively, separated by the subspace $\langle \mathbf{w}, \mathbf{w}_{\mathbf{EC}} \rangle = 0$.

Proof. We perform the change of variable $\mathbf{u} = \sqrt{\mathbf{C}}\mathbf{w}$, so that $\mathbf{u}^t \mathbf{u} = \mathbf{w}^t \mathbf{C} \mathbf{w}$. Notice that $\sqrt{\mathbf{C}}$ is also a symmetric matrix, and that $\mathbf{w} = \sqrt{\mathbf{C}}^{-1} \mathbf{u}$; the system then becomes:

$$\sqrt{\mathbf{C}}^{-1} \dot{\mathbf{u}} = \mathbf{EC} \sqrt{\mathbf{C}}^{-1} \mathbf{u} - (\mathbf{u}^t \sqrt{\mathbf{C}}^{-1} \mathbf{C} \sqrt{\mathbf{C}}^{-1} \mathbf{u}) \sqrt{\mathbf{C}}^{-1} \mathbf{u} = \mathbf{E} \sqrt{\mathbf{C}} \mathbf{u} - (\mathbf{u}^t \mathbf{u}) \sqrt{\mathbf{C}}^{-1} \mathbf{u}$$

or equivalently:

$$\dot{\mathbf{u}} = \sqrt{\mathbf{CE}} \sqrt{\mathbf{C}} \mathbf{u} - (\mathbf{u}^t \mathbf{u}) \mathbf{u} = \mathbf{A} \mathbf{u} - (\mathbf{u}^t \mathbf{u}) \mathbf{u} \quad (7)$$

where, of course, we defined $\mathbf{A} = \sqrt{\mathbf{CE}} \sqrt{\mathbf{C}}$. Clearly, \mathbf{A} a symmetric matrix, having the same eigenvalues as \mathbf{EC} . More precisely, \mathbf{w} is an eigenvector of \mathbf{EC} with eigenvalue μ iff $\sqrt{\mathbf{C}}\mathbf{w}$ is an eigenvector of \mathbf{A} with eigenvalue μ . Moreover: any two distinct eigenvectors $\mathbf{v} \neq \mathbf{w}$ of \mathbf{EC} are known to be orthogonal, hence any two distinct eigenvectors of \mathbf{A} are orthogonal in the regular Euclidean dot product: $(\sqrt{\mathbf{C}}\mathbf{v})^t (\sqrt{\mathbf{C}}\mathbf{w}) = \mathbf{v}^t \sqrt{\mathbf{C}} \sqrt{\mathbf{C}} \mathbf{w} = \mathbf{v}^t \mathbf{C} \mathbf{w} = 0$.

Consider then \mathbf{v} to be the principal component of \mathbf{A} (i.e., the eigenvector corresponding to its maximal eigenvalue), and let $\mathbf{u} = \mathbf{u}(t)$ be a trajectory of the system (7). We want to observe the evolution in time of the angle between the variable vector \mathbf{u} and the fixed vector \mathbf{v} .

$$\cos \theta = \frac{\langle \mathbf{v}, \mathbf{u} \rangle}{\|\mathbf{v}\| \cdot \|\mathbf{u}\|}$$

We differentiate and obtain:

$$-\|\mathbf{v}\| \sin(\theta) \dot{\theta} = \frac{1}{\|\mathbf{u}\|^2} \left[\langle \mathbf{v}, \dot{\mathbf{u}} \rangle \cdot \|\mathbf{u}\| - \langle \mathbf{v}, \mathbf{u} \rangle \frac{\langle \mathbf{u}, \dot{\mathbf{u}} \rangle}{\|\mathbf{u}\|} \right] = \frac{(\mathbf{v}^t \dot{\mathbf{u}}) \|\mathbf{u}\|^2 - (\mathbf{v}^t \mathbf{u}) (\mathbf{u}^t \dot{\mathbf{u}})}{\|\mathbf{u}\|^3} \quad (8)$$

The numerator of this expression

$$\begin{aligned} h(\mathbf{u}) &= (\mathbf{v}^t \dot{\mathbf{u}}) (\mathbf{u}^t \mathbf{u}) - (\mathbf{v}^t \mathbf{u}) (\mathbf{u}^t \dot{\mathbf{u}}) = (\mathbf{u}^t \mathbf{u}) \left(\mathbf{v}^t [\mathbf{A} \mathbf{u} - (\mathbf{u}^t \mathbf{u}) \mathbf{u}] \right) + (\mathbf{u}^t \mathbf{v}) - (\mathbf{v}^t \mathbf{u}) \left(\mathbf{u}^t - [\mathbf{A} \mathbf{u} - (\mathbf{u}^t \mathbf{u}) \mathbf{u}] \right) \\ &= (\mathbf{u}^t \mathbf{u}) (\mathbf{v}^t \mathbf{A} \mathbf{u}) - (\mathbf{v}^t \mathbf{u}) (\mathbf{u}^t \mathbf{A} \mathbf{u}) \end{aligned}$$

We are interested in the sign of $h(\mathbf{u})$; to make our computations simpler, we can diagonalize \mathbf{A} in a basis of orthogonal eigenvectors $\mathbf{A} = \mathbf{P}^t \mathbf{D} \mathbf{P}$, where \mathbf{D} is the diagonal matrix of eigenvalues and \mathbf{P} is an orthogonal matrix whose columns are the eigenvectors. Then:

$$\begin{aligned} h(\mathbf{u}) &= [(\mathbf{P} \mathbf{u}^t) (\mathbf{P} \mathbf{u})] [(\mathbf{P} \mathbf{v}^t) \mathbf{D} (\mathbf{P} \mathbf{u})] - [(\mathbf{P} \mathbf{v}^t) (\mathbf{P} \mathbf{u})] [(\mathbf{P} \mathbf{u}^t) \mathbf{D} (\mathbf{P} \mathbf{u})] \\ &= (\mathbf{z}^t \mathbf{z}) (\mathbf{y}^t \mathbf{D} \mathbf{z}) - (\mathbf{y}^t \mathbf{z}) (\mathbf{z}^t \mathbf{D} \mathbf{z}) \end{aligned}$$

where $\mathbf{y} = \mathbf{P} \mathbf{v}$ and $\mathbf{z} = \mathbf{P} \mathbf{u}$, so that $\mathbf{D} \mathbf{y} = \mathbf{D} \mathbf{P} \mathbf{v} = \lambda_1 \mathbf{y}$ (where $\lambda_1 > \lambda_2 \geq \dots \geq \lambda_n$ is the largest eigenvalue of \mathbf{EC} , assumed to have multiplicity one. Hence:

$$\begin{aligned} h(\mathbf{u}) &= (\mathbf{z}^t \mathbf{z}) (\mathbf{y}^t \mathbf{D} \mathbf{z}) - (\mathbf{y}^t \mathbf{z}) (\mathbf{z}^t \mathbf{D} \mathbf{z}) = (\mathbf{z}^t \mathbf{z}) \lambda_1 (\mathbf{y}^t \mathbf{z}) - (\mathbf{y}^t \mathbf{z}) (\mathbf{z}^t \mathbf{D} \mathbf{z}) = (\mathbf{y}^t \mathbf{z}) [\lambda_1 (\mathbf{y}^t \mathbf{z}) - \mathbf{z}^t \mathbf{D} \mathbf{z}] \\ &= (\mathbf{y}^t \mathbf{z}) \left[\lambda_1 \sum z_j^2 - \sum \lambda_j z_j^2 \right] = (\mathbf{y}^t \mathbf{z}) \left[\sum (\lambda_1 - \lambda_j) z_j^2 \right] \end{aligned}$$

Hence, if $\mathbf{y}^t \mathbf{z} > 0$, then $h(\mathbf{u}) > 0$. In other words: if $\mathbf{v}^t \mathbf{u} > 0$ then $-\|\mathbf{v}\| \sin(\theta) \dot{\theta} > 0$, hence that $\dot{\theta} < 0$. For our original system, this means that any trajectory starting at a \mathbf{w} with $\langle \mathbf{w}, \mathbf{w}_{\mathbf{EC}} \rangle > 0$ converges in time towards the principal eigenvector $\mathbf{w}_{\mathbf{EC}}$ of the matrix \mathbf{EC} . □

Appendix 3. Sensitivity analysis

This is a technical section, in which we calculate how the invariant directions $z_{1,2}$ change when varying q .

Remark. In order to simplify further computations, we rewrite:

$$\begin{aligned} z_{1,2} &= \frac{-q\delta \pm \sqrt{\Delta}}{2\beta} = \frac{-q\delta \pm \sqrt{q^2\delta^2 + 4\beta[\beta + (1-q)\delta]}}{2\beta} \\ &= -\frac{1}{2} \left(\frac{q\delta}{\beta} \right) \pm \frac{1}{2} \text{sign}(\beta) \sqrt{\left(\frac{q\delta}{\beta} \right)^2 + \frac{4[\beta + (1-q)\delta]}{\beta}} \end{aligned}$$

Call $\gamma = \frac{q\delta}{\beta} = \frac{q\delta}{cq + (1-q)v}$. Then $\frac{(1-q)\delta}{\beta} = \frac{\delta - c\gamma}{v}$, and hence:

$$z_{1,2} = -\frac{1}{2}\gamma \pm \frac{1}{2}\text{sign}(\beta)\sqrt{\eta}$$

where

$$\eta = \frac{\Delta}{\beta^2} = \gamma^2 + 4 \left[1 + \frac{\delta - c\gamma}{v} \right]$$

Then we can use the chain rule to express $\frac{dz_{1,2}}{dq} = \frac{dz_{1,2}}{d\gamma} \cdot \frac{d\gamma}{dq}$.

Lemma. The derivative $\frac{d\gamma}{dq} = \frac{\delta v}{\beta^2}$. Also, for $q \in (1/2, q^*) \cup (q^*, 1]$ (i.e., where $\beta \neq 0$), we have:

$$\frac{dz_{1,2}}{d\gamma} = \frac{1}{2} \left[\frac{\pm \left(q\delta - \frac{2c\beta}{v} \right)}{\sqrt{\Delta}} - 1 \right] < 0$$

Proof. $\frac{d\gamma}{dq} = \frac{d}{dq} \left(\frac{q\delta}{\beta} \right) = \frac{\delta[\beta - q\dot{\beta}]}{\beta^2} = \frac{\delta[\beta - q(c-v)]}{\beta^2} = \frac{\delta v}{\beta^2}$

For $q \in (1/2, q^*) \cup (q^*, 1]$, we also have directly that:

$$\begin{aligned} \frac{dz_{1,2}}{d\gamma} &= -\frac{1}{2} \pm \frac{1}{2} \frac{\text{sign}(\beta) \frac{d\eta}{d\gamma}}{\sqrt{\eta}} = -\frac{1}{2} \pm \frac{1}{2} \text{sign}(\beta) \frac{\gamma - \frac{2c}{v}}{\sqrt{\gamma^2 + 4 \left[1 + \frac{\delta - c\gamma}{v} \right]}} \end{aligned} \quad (9)$$

$$= \frac{1}{2} \left[\pm \frac{\beta \left(\gamma - \frac{2c}{v} \right)}{\sqrt{\Delta}/|\beta|} - 1 \right] = \frac{1}{2} \left[\frac{\pm \left(q\delta - \frac{2c\beta}{v} \right)}{\sqrt{\Delta}} - 1 \right] \quad (10)$$

Since $\left(\gamma^2 + 4 \left[1 + \frac{\delta - c\gamma}{v} \right] \right) - \left(\gamma - \frac{2c}{v} \right)^2 = \frac{4[v(v+\delta) - c^2]}{v^2} > 0$, it follows that :

$$\sqrt{\gamma^2 + 4 \left[1 + \frac{\delta - c\gamma}{v} \right]} > \left| \gamma - \frac{2c}{v} \right| \geq \pm \left(\gamma - \frac{2c}{v} \right)$$

and hence

$$\frac{\left| \text{sign} \beta \left(\gamma - \frac{2c}{v} \right) \right|}{\sqrt{\gamma^2 + 4 \left[1 + \frac{\delta - c\gamma}{v} \right]}} < 1$$

It immediately follows from (4) that $\frac{dz_{1,2}}{d\gamma} < 0$. □

Corollary. *The slope of the invariant lines changes with respect to q according to:*

$$\frac{dz_{1,2}}{dq} = \frac{\delta v}{2\beta^2} \left[\frac{\pm \left(q\delta - \frac{2c\beta}{v} \right)}{\sqrt{\Delta}} - 1 \right]$$

hence $\text{sign} \left(\frac{dz_{1,2}}{dq} \right) = -\text{sign}(\delta)$ for $q \in (1/2, q^*) \cup (q^*, 1]$.

Proof. The conclusion follows directly from the chain rule that $\frac{dz_{1,2}}{dq} = \frac{\delta v}{\beta^2} \cdot \frac{dz_{1,2}}{d\gamma}$. □

At this stage, we can distinguish two cases: $\delta < 0$ and $\delta > 0$. We analyze in detail the case $\delta > 0$. The other is very similar (although not symmetric about $\delta = 0$), and we will only state the results, and show some graphic illustrations.

Proposition 2.6. *If $\delta > 0$, then $\frac{dz_{1,2}}{dq} < 0$; hence both $z_{1,2}$ are decreasing as $q \in (1/2, q^*) \cup (q^*, 1]$. Furthermore, the monotonicity, asymptotes and end behavior of the functions $z_{1,2}(q)$ are sketched in the following table:*

q	$1/2$		q^*		1
z_1	1	\searrow	$\frac{1 - q^*}{q^*}$	\searrow	$\frac{\delta - \sqrt{4c^2 + \delta^2}}{-2c}$
z_2	$-1 - \frac{\delta}{2(v+c)}$	\searrow	$-\infty \infty$	\searrow	$\frac{\delta + \sqrt{4c^2 + \delta^2}}{-2c}$

Remark 1. In the system's phase plane, this corresponds to a continuous clockwise rotation of the two invariant lines (the vertical asymptote at q^* corresponds to the z_2 line going through a the vertical position). A phase-plane sketch of this process is shown in Figure 2 in the main text, and the graphs of the actual functions $z_{1,2}(q)$ and of their derivatives $dz_{1,2}/dq$, for some fixed values of the parameters $v, c, \delta > 0$, are shown in the Figure below.

Remark 2. Clearly from the table, the angular position of the two equilibria at $q = q^*$ does not depend on the bias δ . It can be easily shown that the norm of these points is also independent on δ . For example, the norm of the stable equilibrium is:

$$\|\mathbf{w}\|^2 = \frac{\mu_1(z_1^2 + 1)}{vz_1^2 + 2cz_1 + (v + \delta)} = \frac{(1 - q^*)c + q^*(v + \delta)}{v(1 - q^*)^2 + 2cq^*(1 - q^*) + (v + \delta)q^{*2}} \cdot q^{*2}(z_1^2 + 1) \quad (11)$$

$$= \frac{(1 - q^*)c + q^*(v + \delta)}{q^*[(1 - q^*)c + q^*(v + \delta)]} \cdot q^{*2}(z_1^2 + 1) = \frac{1 - 2q^* + 2q^{*2}}{q^*} \quad (12)$$

Hence the position of the two equilibria at critical quality does is the same for all bias values $\delta > 0$.

Proof. The monotonicity follows from the Corollary. The limit values follow from direct computation. For example:

$$\lim_{q \rightarrow q_+^*} z_2 = \lim_{q \rightarrow q_+^*} \frac{-q\delta - \sqrt{\Delta}}{2\beta} = \lim_{q \rightarrow q_+^*} \frac{-q^*\delta}{0^-} = +\infty$$

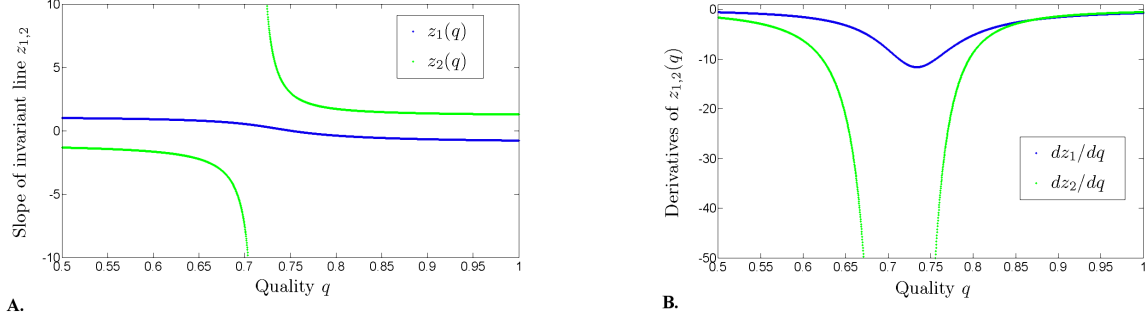


Figure:

Slopes of invariant lines (A) and their change as q is varied (B). In both panels, the other parameters values were fixed to $v = 1$, $c = -0.4$ and $\delta = 0.2$. Notice that, in accordance with Proposition 2.6, z_1 and its derivative dz_1/dq are continuous (blue curves) on $[1/2, 1]$, while z_2 and its derivative dz_2/dq (green curves) have vertical asymptotes at $q = q^* = \frac{v}{v-c} \sim 0.71$.

$$\lim_{q \rightarrow q_-^*} z_2 = \lim_{q \rightarrow q_-^*} \frac{-q^* \delta}{0^+} = -\infty$$

$$\lim_{q \rightarrow q^*} z_1 = \lim_{q \rightarrow q^*} \frac{-q\delta + \sqrt{\Delta}}{2\beta} \cdot \frac{-q\delta - \sqrt{\Delta}}{-q\delta - \sqrt{\Delta}} = \lim_{q \rightarrow q^*} \frac{\beta + (1-q)\delta}{q\delta + \sqrt{\Delta}} = \frac{1-q^*}{q^*} \quad \square$$

Remark 3. If $\delta < 0$, then $\frac{dz_{1,2}}{dq} > 0$ and hence both z_1 and z_2 are increasing as $q \in (1/2, q^*) \cup (q^*, 1]$. In the system's phase plane, this corresponds to a continuous counter-clockwise rotation of the two invariant lines.

Proposition 2.7. For $\delta > 0$, the angle $\theta_{1,2} \in [-\pi/2, \pi/2]$ between each invariant line and the w_1 abscissa is decreasing with respect to the parameter q . Moreover, the angular rate of change is finite, at all $q \in (1/2, 1]$.

q	$\frac{1}{2}$	q^*	1
$d\theta_1/dq$	(-)	$-\frac{\det(C)}{v\delta(1-2q^*+2q^2)}$	(-)
$d\theta_2/dq$	(-)	$-\frac{v}{q^2\delta}$	(-)

Proof. The relation between the slope z and the actual angle θ is given by: $z = \tan \theta$ (we will avoid indices wherever there is no danger of confusion). Hence, for $q \in (1/2, q^*) \cup (q^*, 1]$, we have:

$$\cos^2(\theta) \cdot \frac{d\theta}{d\gamma} \implies \frac{d\theta}{d\gamma} = \frac{dz}{d\gamma} \cdot \frac{dz}{dz} \cdot \frac{1}{z^2 + 1}$$

So

$$\frac{d\theta}{dq} = \frac{d\gamma}{dq} \cdot \frac{d\theta}{d\gamma} = \frac{\delta v}{2\beta^2} \cdot \frac{dz}{d\gamma} \cdot \frac{1}{z^2 + 1} \quad (13)$$

hence $\text{sign}\left(\frac{d\theta}{dq}\right) = -\text{sign}(\delta)$, for all $q \in (1/2, q^*) \cup (q^*, 1]$.

We yet have to check that the rate of change $\frac{d\gamma}{dq}$ remains finite (i.e., does not blow up to $-\infty$) as $q \rightarrow q^*$. Elaborating on (6) we have, for $q \in (1/2, q^*) \cup (q^*, 1]$:

$$\begin{aligned} \lim_{q \rightarrow q^*} \frac{d\theta_2}{dq} &= \lim_{q \rightarrow q^*} \frac{\delta v}{2\beta^2} \cdot \frac{-(q\delta - \frac{2c\beta}{v}) - \sqrt{\Delta}}{\sqrt{\Delta}} \cdot \frac{4\beta^2}{\Delta + 4\beta^2 - 2q\delta\sqrt{\Delta}} \\ &= \frac{2\delta v}{q^*\delta} \cdot \frac{-q^*\delta - q^*\delta}{2q^{*2}\delta^2 + 2q^{*2}\delta^{*2}} = -\frac{v}{q^{*2}\delta} \end{aligned} \quad (14)$$

We also notice that

$$\lim_{q \rightarrow q^*} z_1 = \frac{1 - q^*}{q^*} \implies \lim_{q \rightarrow q^*} (z_1^2 + 1) = \frac{1 - 2q^* + 2q^{*2}}{q^{*2}} \quad (15)$$

and that

$$\begin{aligned} \lim_{q \rightarrow q^*} \frac{dz_1}{dq} &= \lim_{q \rightarrow q^*} \frac{\delta v}{2\beta^2} \cdot \frac{\left(q\delta - \frac{2c\beta}{v}\right) - \sqrt{\Delta}}{\sqrt{\Delta}} = \lim_{q \rightarrow q^*} \frac{\delta v}{2\beta^2} \cdot \frac{\left(q\delta - \frac{2c\beta}{v}\right) - \sqrt{\Delta}}{\sqrt{\Delta}} \cdot \frac{\left(q\delta - \frac{2c\beta}{v}\right) + \sqrt{\Delta}}{\left(q\delta - \frac{2c\beta}{v}\right) + \sqrt{\Delta}} \\ &= \lim_{q \rightarrow q^*} \frac{\delta v}{2\beta^2 \sqrt{\Delta}} \cdot \frac{\left(q\delta - \frac{2c\beta}{v}\right)^2 - \Delta}{\left(q\delta - \frac{2c\beta}{v}\right) + \sqrt{\Delta}} = \lim_{q \rightarrow q^*} \frac{\delta v}{2\beta^2 \sqrt{\Delta}} \cdot \frac{\frac{-4\beta^2}{v^2}(c^2 - v^2 - v\delta)}{\left(q\delta - \frac{2c\beta}{v}\right) + \sqrt{\Delta}} \\ &= \lim_{q \rightarrow q^*} \frac{\delta v}{2\beta^2 \sqrt{\Delta}} \cdot \frac{-4\beta^2 \det(\mathbf{C})}{v^2} \cdot \frac{1}{\left(q\delta - \frac{2c\beta}{v}\right) + \sqrt{\Delta}} \\ &= \frac{\delta v}{2\beta^2 q^* \delta} \cdot \frac{-4\beta^2 \det(\mathbf{C})}{v^2} \cdot \frac{1}{2q^* \delta} = \frac{-\det(\mathbf{C})}{v\delta q^{*2}} \end{aligned} \quad (16)$$

Combining (9) and (10), we have:

$$\lim_{q \rightarrow q^*} \frac{d\theta_2}{dq} = \frac{dz_1}{dq} \frac{1}{z_1^2 + 1} = \frac{-\det(\mathbf{C})}{v\delta q^{*2}} \cdot \frac{q^{*2}}{1 - 2q^* + 2q^{*2}} = \frac{-\det(\mathbf{C})}{v\delta(1 - 2q^* + 2q^{*2})} \quad (17)$$

□

Appendix 4. Description of the ellipse attractor

For unbiased inputs $\delta = 0$ and critical quality $q = q^*$, \mathbf{EC} has a double eigenvalue $\mu = v + c = (2q^* - 1)(v - c)$. The eigenspace of \mathbf{EC} is \mathbb{R}^2 , hence each direction (i.e. slope $z = \tan \theta \in [-\infty, +\infty]$) produces two equilibria, normalized as follows:

$$\begin{aligned}\|\mathbf{w}\|^2 &= \frac{\mu(z^2 + 1)}{vz^2 + 2cz + v} = \frac{(v + c) [\tan^2 \theta + 1]}{v \tan^2 \theta + 2c \tan \theta + v} = \frac{v + c}{v \sin^2 \theta + 2c \sin \theta \cos \theta + v \cos^2 \theta} \\ &= \frac{v + c}{v + c \sin(2\theta)}\end{aligned}\tag{18}$$

We show that this is the polar equation of an ellipse with foci along the first diagonal $\theta = \pi/4$. Indeed, under a clockwise rotation by $-\pi/4$, equation (13) becomes:

$$\begin{aligned}\rho^2 &= \frac{v + c}{v + c \sin(2[\theta - \frac{\pi}{4}])} = \frac{v + c}{v + c \cos(2\theta)} = \frac{v + c}{v [\cos^2 \theta + \sin^2 \theta] + c [\cos^2 \theta - \sin^2 \theta]} \\ &= \frac{v + c}{(v + c) \cos^2 \theta + (v - c) \sin^2 \theta} = \frac{v + c}{\sqrt{v^2 - c^2}} \cdot \frac{\sqrt{v^2 - c^2}}{(v + c) \cos^2 \theta + (v - c) \sin^2 \theta}\end{aligned}\tag{19}$$

In polar coordinates, this is the equation of an ellipse

$$\rho^2 = \frac{a^2 b^2}{a^2 \cos^2 \theta + b^2 \sin^2 \theta}$$

with radial coordinate $\rho = \|\mathbf{w}\| \frac{\sqrt{v^2 - c^2}}{v + c}$ and angular coordinate θ , semi-major radius $a = \sqrt{v + c}$ and semi-minor radius $b = \sqrt{v - c}$.